Ioannis Kompatsiaris · Jonathan Cave
Anna Satsiou · Georg Carle
Antonella Passani · Efstratios Kontopoulos
Sotiris Diplaris · Donald McMillan (Eds.)

LNCS 10673

Internet Science

4th International Conference, INSCI 2017
Thessaloniki, Greece, November 22–24, 2017
Proceedings

Springer

Lecture Notes in Computer Science 10673

Commenced Publication in 1973
Founding and Former Series Editors:
Gerhard Goos, Juris Hartmanis, and Jan van Leeuwen

More information about this series at http://www.springer.com/series/7409

Ioannis Kompatsiaris · Jonathan Cave
Anna Satsiou · Georg Carle
Antonella Passani · Efstratios Kontopoulos
Sotiris Diplaris · Donald McMillan (Eds.)

Internet Science

4th International Conference, INSCI 2017
Thessaloniki, Greece, November 22–24, 2017
Proceedings

Springer

Editors
Ioannis Kompatsiaris
CERTH
Information Technologies Institute
Thessaloniki
Greece

Jonathan Cave
University of Warwick
Coventry
UK

Anna Satsiou
CERTH
Information Technologies Institute
Thessaloniki
Greece

Georg Carle
Technical University Munich
Garching
Germany

Antonella Passani
T6 Ecosystems
Rome
Italy

Efstratios Kontopoulos
CERTH
Information Technologies Institute
Thessaloniki
Greece

Sotiris Diplaris
CERTH
Information Technologies Institute
Thessaloniki
Greece

Donald McMillan
Stockholm University
Stockholm
Sweden

ISSN 0302-9743 ISSN 1611-3349 (electronic)
Lecture Notes in Computer Science
ISBN 978-3-319-70283-4 ISBN 978-3-319-70284-1 (eBook)
https://doi.org/10.1007/978-3-319-70284-1

Library of Congress Control Number: 2017957858

LNCS Sublibrary: SL3 – Information Systems and Applications, incl. Internet/Web, and HCI

Printed on acid-free paper

This Springer imprint is published by Springer Nature
The registered company is Springer International Publishing AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

Preface

This volume contains the papers presented at INSCI 2017, the 4th International Conference on Internet Science, held during November 22–24, 2017, in Thessaloniki, Greece.

This year the conference featured work that used a range of technical and social approaches to shed light on the path toward a more open, interoperable, human-centric, democratic, and trustworthy Internet. This provides a glimpse of an Internet more conducive to social innovation and services, with more intelligence, and with greater involvement and participation across the board.

In order to provide this insight, INSCI 2017 brought together researchers from around the world and across disciplinary boundaries. Together we are able to explore the sociotechnical nature of the Internet through the lenses of computer science, sociology, art, mathematics, physics, complex systems analysis, psychology, economics, law, political sciences, and more. The interdisciplinary nature of Internet science as a field offers the opportunity to develop a holistic understanding of the Internet, from the technical innovations, through their policy implications, to the combined impacts on society.

This conference was built on the success of the First and Second International Conferences on Internet Science, which both took place in Brussels, Belgium, organized by the FP7 European Network of Excellence in Internet Science (EINS) project, with the support of the European Commission. In its third year, the conference moved to Florence, Italy, with a theme in line with, and support from, the Collective Awareness Platforms for Sustainability and Social Innovation (CAPS) initiative. A continued interest in sustainability and social innovation can also be seen reflected in many of the submissions received this year.

In its fourth year, the theme of the International Conference on Internet Science was "Technical and Social Approaches for a Human-Centric, Democratic, and Trustworthy Next-Generation Internet." This theme aimed to further explore the goals of the Next-Generation Internet Initiative, and to support the overarching goal of the field of Internet Science — to provide the sociotechnical means to exploit the full potential of the Internet for society and the economy. The fourth edition of the conference was organized with the support of the EU CAPS and Next-Generation Internet (NGI) initiatives.

The papers presented in this volume fall into five thematic topics that align with the aforementioned goals:

Next-Generation Community Engagement
Online Policy, Politics, and Co-creation
Understanding and Empowering Digital Citizens
Data-Driven Research and Design
Social Media and Online Interaction

The organizers of the 4th International Conference on Internet Science wish to thank their organizing partners (The Centre for Research and Technology Hellas, University of Warwick, Technical University of Munich, T6 Ecosystems, Nokia Bell Labs, University of Stockholm, University of Southampton, University of Florence, University of Bologna, Winchester University, University of Athens, Centre for Research on Direct democracy) for their help and support.

October 2017

Ioannis Kompatsiaris
Jonathan Cave
Anna Satsiou
Georg Carle
Antonella Passani
Efstratios Kontopoulos
Sotiris Diplaris
Donald McMillan

Organization

Program Committee

General Chairs

Ioannis Kompatsiaris	Centre for Research and Technology Hellas, Greece
Jonathan Cave	University of Warwick, UK

Technical Program Committee Chairs

Georg Carle	Technical University of Munich, Germany
Anna Satsiou	Centre for Research and Technology Hellas, Greece
Antonella Passani	T6 Ecosystems, Italy

Workshops and Sessions Chairs

Sotiris Diplaris	Centre for Research and Technology Hellas, Greece
Vishal Kapadia	Wikirate Project e.V., Germany

Demos and Experimental Showcases Chairs

Dimitri Papadimitiou	Nokia Bell Labs, Belgium
Efstratios Kontopoulos	Centre for Research and Technology Hellas, Greece

Proceedings Chair

Donald McMillan	Stockholm University, Sweden

Sponsorships Chair

Theodora Tsikrika	Centre for Research and Technology Hellas, Greece

Local Arrangements Chair

Maria Papadopoulou	Centre for Research and Technology Hellas, Greece

Website Chair

Katerina Katmada	Centre for Research and Technology Hellas, Greece

Steering Committee

Thanassis Tiropanis	University of Southampton, UK
Laura Sartori	University of Bologna, Italy
Yanina Welp	Centre for Research on Direct Democracy, Switzerland
Ioannis Stavrakakis	University of Athens, Greece
Franco Bagnoli	University of Florence, Italy

Olessia Koltsova	St. Petersburg School of Social Sciences and Humanities, Russia
Caroline Wilson	University of Winchester, UK

Technical Program Committee

Panayotis Antoniadis	Nethood, Switzerland
Dimitrios Bekiaridis-Moschou	Greek Association of Alzheimer's Disease and Related Disorders, Greece
Antonis Bikakis	University College London, UK
Theologos Dergiades	University of Macedonia, Greece
Karmen Guevara	Independent, UK
Clare Hooper	OpenRoad Communications, Canada
Georgios Iosifidis	Trinity College Dublin, Ireland
Konstantinos Kafetsios	University of Crete, Greece
Merkouris Karaliopoulos	Athens University of Economics and Business, Greece
Eirini Karapistoli	Democritus University of Thrace, Greece
Ioannis Katakis	National and Kapodistrian University of Athens, Greece
Katerina Katmada	Centre for Research and Technology Hellas, Greece
Iordanis Koutsopoulos	Athens University of Economics and Business, Greece
Nikolaos Lagos	Naver Labs Europe, France
Ioulietta Lazarou	Centre for Research and Technology Hellas, Greece
Alessandro Mantelero	Politecnico di Torino (Nexa Center), Italy
Chris Marsden	University of Sussex, UK
Meryem Marzouki	LIP6 - CNRS and UPMC Sorbonne Universités – Paris, France
Richard Mills	University of Cambridge, UK
Theodoros Moysiadis	Centre for Research and Technology Hellas/Institute of Applied Biosciences, Greece
Hermann De Meer	University of Passau, Germany
Federico Morando	Nexa Centre for Internet and Society, Italy
Heiko Niedermayer	Technical University of Munich, Germany
Spiros Nikolopoulos	Centre for Research and Technology Hellas, Greece
Nasrine Olson	University of Borås, Sweden
Manos Panaousis	University of Brighton, UK
Symeon Papadopoulos	Centre for Research and Technology Hellas, Greece
Thomas Papadopoulos	University of Cyprus, Greece
Ioannis Partalas	Viseo Group, France
Albert Meroño Peñuela	Vrije Universiteit Amsterdam, The Netherlands
Alessandra Prampolini	T6 Ecosystems, Italy
Mark Rouncefield	Lancaster University, UK
Simona De Rosa	T6 Ecosystems, Italy
Kave Salamatian	University of Savoie, France
Andrea Scharnhorst	Royal Netherland Academy for Arts and Science, The Netherlands

Anthoula Tsolaki	Aristotle University of Thessaloniki and Neurology Clinic G.H. Agios Pavlos, Thessaloniki, Greece
Ziga Turk	University of Ljubljana, FGG, Slovenia
Stefanos Vrochidis	Centre for Research and Technology Hellas, Greece
Jean-Yves Vion-Dury	Naver Labs Europe, France
Simon Waddington	King's College London, UK
Peter Wittek	Institute of Photonic Sciences (ICFO), Spain
Bendert Zevenbergen	Princeton University, USA

Contents

Next Generation Community Engagement

WeMake: A Framework for Letting Students Create Tangible, Embedded and Embodied Environments for Their Own STEAM Learning

Anastasios Karakostas[1(✉)], George Palaigeorgiou[2], and Yiannis Kompatsiaris[1]

[1] Centre for Research and Technology Hellas, Thessaloniki, Greece
{akarakos,ikom}@iti.gr
[2] Department of Primary Education, University of Western Macedonia Florina, Florina, Greece
gpalegeo@uowm.gr

Abstract. This paper presents the principles and the design of the WeMake framework. The goal of the WeMake framework is twofold: firstly, to create an interdisciplinary team of experts that together with students/teachers and a new participatory design methodology adapted to embodied interactions will develop low cost and easily reconstructable embodied interaction environments for STEAM domains; and secondly to invite students, teachers and schools across the world to build, exploit, share and assess their own versions of these embodied learning environments. The ultimate goal is to create an infrastructure that will motivate all stakeholders (from researchers to students) and maintain a perpetual cycle of embodied STEAM learning environment proposals and their deployment in the educational practice.

Keywords: Embodied learning · STEAM · Mixed reality environments · Wearables · Tangibles · Participatory design

1 Introduction

Embodied cognition theory and several relevant frameworks suggest that thinking and acting (or else mind and body) are intertwined in nature and that tangible engagement with objects or exploring spaces affects the way we think about them and vice-versa. Grounded Cognition describes that mental representations are grounded in motor areas of the cortex and that the perceptual and motor states acquired through experience are reactivated through simulation when knowledge is needed [9]. Similarly, the Embodied Metaphor theory suggests that abstract concepts and conceptual metaphors are based on image schemas that derive from physical actions [3]. And many more theoretical frameworks propose that full-body interaction has the potential to support learning by involving users at different levels such as sensorimotor experience, cognitive aspects and affective factors; the physical world seems to underpin one's internal mental

I. Kompatsiaris et al. (Eds.): INSCI 2017, LNCS 10673, pp. 3–18, 2017.
https://doi.org/10.1007/978-3-319-70284-1_1

representations [35]. The embodied framework of learning seems a promising one especially for STEAM disciplines since in order to understand STEAM content students must adjust their unmediated perceptions and actions with the mediated structures of disciplinary practice. The design rationale is that having learners act out and physicalize the systems processes, relationships, etc., will create conceptual anchors from which new knowledge can be built [32].

New interaction technologies can prove an excellent guide for students to perform physical actions that serve as "conceptual leverage" [32]. Under the umbrella of terms like embodied interaction, full-body interaction, motion-based interaction, gesture-based interaction, tangible interaction, bodily interaction, and kinesthetic interaction, several interactive learning environments based on novel interaction modalities have been developed. Following similar theoretical underpinnings, these interactive environments try to facilitate an embodied experience of a certain concept, to represent an abstract concept as a concrete instance or operationalize actions as means to express specific content, or try to use space as a semiotic resource or even try to become embodied metaphors. The new mediated environments seem to increase learner engagement since body-based experiences are more perceptually immersive and learners may feel that they are in a more authentic and meaningful educational space [19].

However, designing learning activities that promote conditions for embodied learning is an emerging and not yet systematized area of research. Some studies have indicated that binding knowledge with physical interactions has a strong effect on learning but nevertheless it is undoubtedly not the case, that any body movement will stimulate learning. Recent studies of embodied learning have focused on the notion of congruency ([1]): the condition under which movements or body positioning map are related to a particular conceptual domain. Further research is required for researchers to learn how to cue the body to enact certain actions and create physical representations that facilitate conceptual understanding.

Existing research on embodied learning technologies has been disparate, driven largely by specific technical innovations and constraints, and often lacking a clear focus on establishing their efficacy in educational contexts. The design space for ideating relevant interactive learning environments is chaotic and most of the proposed environments are polarized either towards the technological or the pedagogical view on the basis of the team members that produce each learning environment. How can we combine all these different opportunities, the different theoretical frameworks and the various technological enablers, the various scientific perspectives, the experience in each separate learning domain in order to construct effective, efficient and enjoyable embodied interactive environments? It is more than evident that designing systems that enable physically immersive and responsive pedagogical activities requires interdisciplinary collaboration between a plethora of experts with different backgrounds, and not only computer experts and learning experts. For example, practices of embodied doing and thinking through action are central to the arts and related to STEAM education. In arts, intelligent practice subtly draws upon and orders materiality – with gestures, breathing, and the artful manipulation of instruments, tools and materials, each with its own qualities. The Arts can offer a lot to embodied STEM education through different perspectives.

Another significant drawback of current embodied interactive technologies is that their mass deployment in schools seems infeasible since they are usually developed only for research purposes, are customized with expensive technologies, do not follow the curriculum while teachers also do not have confidence in using them. Interestingly, this embodied technological approach to STEAM happens to coincide with the recent enthusiasm for the "maker movement" [24]. Low cost rapid prototyping together with the uprising trend of arts and crafts fairs, tinkering and inventing, can enable students to recreate such embodied interactive technologies and concurrently empower creativity and problem solving, brainstorming, sustained perseverance and a lot more. This approach also provides the motivational aspect often missing in traditional STEAM curricula that sometimes serves to disengage students from STEM subjects taught in formal educational settings. Would it be possible for students and teachers to construct their own STEAM learning tools?

In this manuscript, we will present the framework WeMake which proposes a series of steps in order to combine different opportunities, the different theoretical frameworks and the various technological and pedagogical enablers in order to trigger students and teachers to construct their own STEAM learning tools.

2 WeMake: Students Create Tangible, Embedded and Embodied Environments for Their Own STEAM Learning

WeMake is based on two pillars: firstly, it underlines the need to create an interdisciplinary team of experts that together with students/teachers and a new participatory design methodology adapted to embodied interactions, will develop low cost and easily reconstructable embodied interaction environments for STEAM learning domains; and secondly, to invite students, teachers and schools across the world to build, exploit, share and assess their own versions of these embodied learning environments. The ultimate goal is to create an infrastructure that will motivate all stakeholders (from researchers to students) and maintain a perpetual cycle of embodied STEAM learning environment proposals and their deployment in the educational practice beyond the duration of the project.

There are 2 prerequisites:

- The creation of interdisciplinary and multidisciplinary team of experts that will synthesize creatively for the first time an extensive range of different perspectives of interactive embodied learning in STEAM. Expertise is needed in the learning domain (STEAM learning domains), child-computer interaction, cognitive psychology, educational technology, interactive installations, arts, prototyping and participatory design.

- The development of a new participatory design methodology that will leverage this symmetry of ignorance and help the participants negotiate their different perspectives of embodied interactions and also enable them to develop common externalizations of embodied performances that will mediate their collaboration.

In order to explore systematically the design space of interactive embodied learning, the design team will focus on three widely exploited types of interactive embodied learning environments for STEAM:

– *mixed reality spaces* where physical and digital objects co-exist and where the body is situated in the environment and interacts with it. Mixed reality environments offer a large vivid and immersive audiovisual interface for eliciting body activity (interactive floors, interactive walls etc.).
– *tangibles*, digital manipulatives that become objects to think and act with (building bricks and balls, interactive physical objects, tools, etc.).
– *e-textiles* are wearables at the intersection of physical, digital and artful media and which will help students to "embody" new conceptualizations and representations.

The three different interfaces have been selected as they provide multiple means of engagement, multiple means of representation, and multiple means of action and expression.

The design team will offer some first iterations of technologically supported embodied learning (in Maths for example the design team may work on spatial representations of number magnitude and basic arithmetic operations while in Science, the design interactive object may be related to gravity or light).

However, the design products of the interdisciplinary team must comply with two design constraints

- they must be low-cost by using mainstream prototyping hardware and software
- they should be easily reconstructable by students and teachers. In order to be highly populated, the products should be accompanied by construction guides and learning scenarios in order to help students to create them and teachers to embrace them.

Afterwards, all the products will be shared through an online STEAM makerspace platform which will enable researchers, teachers and students to share designs with each other and more specifically to

- to present new embodied interactive learning environments
- to provide learning scenarios and instructional guides for the proposed learning environments
- to present personal variations and improvisations on the proposed learning environments
- to view anonymous learning usage data for each learning environment.

The online STEAM makerspace will promote the maker culture which provides an effective mechanism for K12 students, particularly women and multicultural students, in pursuing STEAM fields.

3 Three Types of Embodied Interaction on Focus

3.1 Embodied STEAM Learning and Mixed Reality Spaces

Mixed Reality (MR) environments merge the digital with the physical and offer a vivid and immersive audiovisual interface for eliciting body activity. In these environments, authentic and expressive physical activity can be augmented with digital displays that emphasize the metaphor and tools for feedback and reflection [30]. Mixed reality technologies allow students to become part of the system they are trying to familiarize with, and give them the advantage of the insider who can monitor and evaluate the mechanisms and relationships that define the domain [30]. In the review of [35] for learning abstract concepts through Full-Body Interaction, the authors identify seventeen learning projects where the physical interface is based on the use of vertical screens or wall projections, requiring users to interact in the space in front of the display and twelve projects which are based on floor projections, which allow the user to move around the periphery of or directly on the visual output. Most of them concern Maths and Science projects. For example, in [21] elementary school students participated with their bodies in simulations of force and motion that allowed them to successfully connect their physical activity with formal representations of these phenomena. Using the Kinect sensor, [33] also pursued the idea of students walking along an up to 3 m long line to explore continuous number line. The MEteor simulation game [31], identify middle school students with asteroids and lets them build insights about how objects move in space. Students are provoked to predict the route of an asteroid as it comes into the vicinity of gravitational forces (e.g., nearby planets). Situated Multimedia Arts Learning Lab (SMALLab) is a 15-foot-by-15-foot space with an interactive floor display that supports simulations and learning modules ranging from language arts and science to technology, engineering, and mathematics education [25, 52].

Up to this point, there has been little research examining outcomes for students learning the same content in the same simulation environment, differing only in the degree of immersion and physical interaction with the interface. Also, the duration of the interventions in the various studies is limited (from 20 min up to 1 week) while relatively few empirical studies have investigated learning in authentic classroom settings [30]. WeMake framework suggests to focus on large interactive floors and will try to exploit varying degrees of sensorimotor activation, gestural congruency and immersion and also explore the collaborative perspective of mixed reality spaces since it is of paramount importance for realistic classroom settings.

3.2 Embodied STEAM Learning and Tangibles

[26], early on, proposed the idea of the ‘tangible bits’ and argued that tangible user interfaces serve the function of making “digital information directly manipulatable with our hands, and perceptible through our peripheral senses by physically embodying it”. Many researchers have suggested that tangible user interfaces (TUIs) have potential for supporting children’s informal and formal learning and that they are highly suited to the design and development of learning activities because they leverage both familiar

physical artifacts and digital computation. Tangibles are frequently used to teach children abstract concepts, in STEAM [36]. For example, Button Matrix [16] uses coupled tactile, vibration and visual feedback to highlight features of a physical experience with arithmetic concepts and cue reflection on the links between the physical experience and mathematical symbols. Tangible Interactive Microbiology environment [29] offers to students an interface with microbiological living cells and tries to promote artistic expression and scientific exploration. Fractangi is an interactive tangible number line which functions as a conceptual metaphor for helping students to understand and exploit fractions by acting with their hands [37] while in FingerTrips students learn about geography by travelling across Europe over an augmented 3D map with the use of their fingers [40].

Tangibles are able to offer a natural and immediate form of interaction that is accessible to learners, promote active and hands-on engagement, allow for exploration, expression, discovery and reflection, provide learners with ‘tools to think with’ and offer opportunities for collaborative activity among learners (e.g. [4, 44, 45]). However, there is also another stream of research which indicates that “physicality is not important” and rather “their manipulability and meaningfulness make them [manipulatives] educationally effective” [46]. In many situations, children do not transfer performance with physical to symbolic representations of problems. Indeed, it has been suggested that previously identified virtues of physical manipulatives—“learning through concrete and perceptually rich physical practices ”—are not the drivers of learning (e.g., [55] and can even be detrimental to learning (e.g., [48]). However, a recent meta-analysis found that the use of physical manipulatives in math education tends to improve retention, problem solving, and transfer [13]. Additionally, the context of use seems to also have detrimental effects. For example, unconstrained physical manipulation has also been shown to be suboptimal for learning [43, 50] or high interactivity can be overwhelming and may lead to a lower learning performance embedded learning (e.g., [50]). Hence, the design of tangibles still holds great difficulty. Tangibles may differ in terms of the metaphorical relationship between the physical and digital representation. Small representational differences may have great effect on performance differences [22].

In order to achieve the goal of designing efficient and effective learning tangibles, designers and researchers have to bring together specific knowledge about children’s cognitive, physical, emotional, and social skills, the idiosyncratic characteristics and prior experience on each field domain and the opportunities provided of tangibles environments. WeMake suggest to investigate the cueing possibilities of tangible representations by exploiting a range of material properties (e.g., size, shape, texture, temperature, weight), by putting emphasis on multiple output modalities of the designed products and by identifying ways to transfer performance from physical to symbolic representations of problems. Although metaphorical relations are promoted in tangible and embodied interaction research, little is known about how to identify and select embodied metaphors, or how to transform them effectively into interaction models [6]. WeMake also proposes to develop and share design tools, similar to Developmentally Situated Design card set [6] for designing tangibles for STEAM learning.

3.3 Embodied STEAM Learning and E-Textiles

Wearable technologies can incorporate a wide variety of sensors for measuring mechanical information (i.e. position, displacement), acoustic information (i.e. volume, pitch), biological information (i.e. heart rate, temperature), optical information (i.e. refraction, light wave frequency) and environmental information (i.e. temperature, humidity) [8] but they can also offer multimodal output by providing haptic feedback, integrating screens, producing sounds [11]. Wearables have been used in education in several forms, i.e. wearable GPS devices in afterschool clubs [51], augmented reality glasses in educational psychology and organizational behaviour classes [14], wearable video cameras for classroom teachers [47], accelerometer enhanced gloves into immersive and interactive museum simulations [34], e-textiles for representing the body internals [38]. Wearables technologies offer several affordances for learning such as in-situ information, recording and feedback [14], communication and distribution of resources [18]. There is also a shift from individuals using wearable technologies in isolation to more socially oriented uses of data [11]. For example, wearable technologies may empower users to exchange fitness data, play games together in real-time, or even see an event from someone else's viewpoint [54]. Another example of a social oriented use of data is using wearable "thinking tags," with which students were able to explore how diseases spread through a population during an interpersonal interaction [15].

WeMake framework proposes to focus more on e-textiles which are conceived as aesthetically compelling designs, at the intersection of physical, digital and artful media and which will help students to "embody" new conceptualizations and representations. E-textiles address excellently the challenge of designing STEAM activities which must authentically engage participants in both STEM fields and the arts. E-textiles have already demonstrated its capacity in the professional realm to invite and sustain participation from women [12]. Electronic textile (e-textile) toolkits have been successful in empowering users to engage in meaningful design that is creative, self-expressive, and personal and concurrently seem to offer greater transparency into STEM disciplinary content i.e. for exploring circuitry [41], or physiology [38]. E-textiles have a lot to offer to students during their construction since students have to garner expertise in several content areas as well as develop the skill sets to think across traditional disciplinary boundaries. Based on this reasoning MakerShoe [27] was constructed as an e-textile construction kit aiming at engaging young children (K-4) in wearable design.

As most researchers suggest [41], there is a need to work on the efficacy of e-textile integration in schools and home as pathways to learning. As it appears from the previous section, wearables, and more specifically e-textiles have been rarely used in the context of embodied learning and usually without exploiting the output modalities that they can offer. WeMake recommends the wearables to be designed by an extended multidisciplinary team of experts together with students and teachers, with the aim of exploring different ways of creating representations with the body. The design of wearable interfaces should identify the preconditions for delivering open-ended exploration, a high degree of personal expression, and aesthetically compelling possibilities with the e-textiles. The design team has to understand the tradeoffs between guiding a child's attention, simplifying concepts, and allowing for the accurate construction of knowledge. Additionally, it has to envision new materials to support embodied learning with e-textiles that is more open to expressive ways of participating in the classroom.

4 The WeMake Phases

4.1 The Overall Approach

As indicated previously, the WeMake aim is to enhance STEAM educational practice by applying interdisciplinary embodied interaction/learning paradigms and by inviting students and teachers to design, develop and exploit their own mixed reality spaces, learning e-textiles, and learning tangibles.

Its structure has been developed having in mind the following goals:

- To develop and study innovative embodied learning artifacts that address significant STEAM learning obstacles
- To provide playful and effective learning interactions to students by integrating museum-inspired digital installations, wearables and learning tangibles in the daily learning practice inside the school
- To make accessible such learning experiences to all schools by exploiting the new wave of low cost prototyping platforms, by addressing curriculum needs and by providing teachers the necessary deployment tools
- To enable students to become the creators of their learning tangibles, mixed reality spaces e-textiles of their schools
- To motivate creativeness and handcraft mentality and to disseminate the maker culture in schools
- To offer synergies between schools of different areas/countries
- To provide personalized embodied interactions by exploiting the wealth of data collected from different schools
- Offer incentives to researchers, companies, learners and teachers for the continuous design, development, application and research analysis of similar embodied learning environments beyond the projects' lifespan.

In the following Fig. 1, the 3 phases of WeMake are presented.

In the following paragraphs, we will describe the three essential ingredients of WeMake, the online makerplace, the participatory design methodology and the dissemination of the design proposals of the interdisciplinary team.

4.2 Online Makerspace for Fostering Maker Culture and Incorporating Research in the Developing Cycle

The maker movement [20], an extension of Do-It-Yourself (DIY) culture, has largely grown out of environments like Hackerspaces and Makerspaces where hobbyists can create and share their designs. Makers usually are self-organized in communities of practices and borrow and adapt (or remix) each others' designs. Online communities,

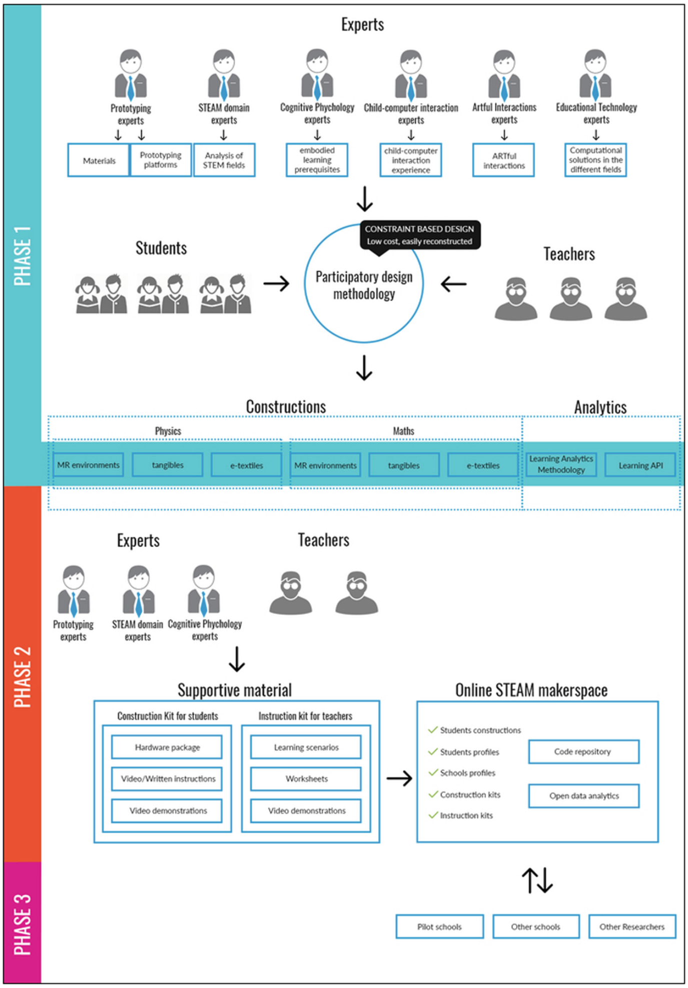

Fig. 1. WeMake phases

such as https://www.thingiverse.com/, https://www.youmagine.com/, http://www.instructables.com/ and https://diy.org/ are an important piece of this ecosystem, allowing geographically distributed makers to share designs with each other [39]. In this new era of web 2.0 collaboration, makers collectively contribute to the development of new designs by iteratively remixing and refining one another's work. Makers' main motivations for contributing to communities include finding inspiration for future projects, learning new concepts, and receiving feedback on projects [28]. Coupling STEAM learning with the makerspace model has the potential to re-center social dimensions of learning to foster supportive collaborative learning communities. The basic characteristics of online maker communities is that of social scaffolding. The dimension of social scaffolding reflects the status in which learners help and inspire one another to persist and solve problems supports an individual's learning [10]. Learners establish social scaffolding by assisting one another while working on the same project or across different projects (i.e. requesting or offering help, inspiring new ideas, offering explanations for a strategy, tool or outcome, striving to understand etc.). Such communities of practices should foster deep engagement of young people, provide opportunities for developing and authoring ideas, and offer the potential for the development of new dispositions, understandings, and directions [53]. Usually online maker communities borrow a lot of elements from physical makerspaces. [42] proposes that physical makerspaces should "provide opportunities to engage with the work and ideas of others; to be supported with tools and assistance to develop and pursue one's own ideas; and to develop and evolve these ideas as direct engagement with materials and phenomena provides feedback, creates constraints, and inspires new thinking and solutions". [10] identified seven central learning practices to collectively characterize recognizable participation in the maker community for STEM-rich tinkering: explore & question, tinker, test, and iterate, hack & repurpose, combine and complexify, customize, seek out resources, and share. The process of externalizing ideas and building on them creates two ideal conditions for learning [41]: (a) externalizing an idea, in words or through an artifact, requires restructuring that idea into different formats (b) creating a physical depiction of an idea and reflecting on that design creates an opportunity for formative feedback.

Introduction of maker culture activities to for embodied STEM learning environments between students, teachers and research is totally new. In WeMake we propose the exploration of the opportunities to create an online maker space where students with an interest in Science, Technology, Engineering, Art and Math (STEAM) can meet up, work on projects, and learn through "Doing-it-Yourself" or "Doing-it-Together". It should also be a place where researchers, teachers and students of all ages can come together and help and learn from one another. WeMake aim is to promote learners to find and pursue a purpose, to exercise their creativity and imagination, and to confront and solve conceptual challenges, within a STEM-rich tinkering context [42], teachers can find valuable info about exploiting embodied STEAM learning environments in and out of the classroom, researchers can propose their own embodied STEAM learning environments together with instructional plans, videos etc. for their deployment and evaluation. WeMake intention is to create a shared repertoire, or a set of shared resources used in the production of new meaning, which includes participatory and reified aspects of practice.

More practically, the online makerplace should enable researchers, teachers and students

- to share the new embodied interactive learning environments (how it works, how to construct)
- to provide learning scenarios and instructional guides for the proposed learning environments (how to deploy in school)
- to present personal variations and improvisations on the proposed learning environments along with the production stages (how do I construct it)

– to view anonymous learning usage data for each learning environment (how it was used, the learning effects)
– to share personal info and discuss, question, comment, rate, favorite all proposed resources

The makerplace should motivate support the user's overall value creation. To reward participants' active involvement, the platform will provide its own badges system and a gamification mechanism as an extrinsic motivational tool. Students, teachers and researchers, by uploading their constructions and instructional proposals, by answering questions, by providing comments and answers to questions, by "liking" other participants' work, will take part in a point accumulation game which will lead to earning badges. The badges should be organized in different perspectives (motivation, pedagogy, and credential) and follow an accreditation system with the use of educational milestones. Such social gamification framework should be tested in real classroom scenarios to obtain empirical data on the actual effects of the gamified processes.

4.3 The Design Methodology

WeMake requires the development of a new participatory design methodology since the interaction in such embodied learning environments is not the only the mean but also the goal. The act of interacting should also be an embodied act for learning; students are asked to act with their bodies, learn and interact with the environment. While a wealth of participatory design methods has been proposed by researchers in the field of Child Computer Interaction, they often are quite open with respect to the design aims, they make excessive demands on the time children can invest, and leave questions regarding the equity of participation: researchers and developers clearly benefit by children's participation, who are asked to provide services on a voluntary basis. The new participatory methodology should be sustainable in the sense that the design activity must provide developmental, educational, self-actualisation, and fun benefits to children in order to keep them engaged at an equitable basis. The new methodology should emerge from tailoring and mixing existing participatory design methods (e.g. cooperative inquiry, future workshop), techniques (e.g. play-back theater, bodystorming) and ideation tools (e.g. ideation decks, technological probes) after making them suitable

(a) for the three types of embodied interaction: mixed reality spaces, tangibles and wearables
(b) for facilitating multi-stakeholders preparation and collaboration (from students to researchers),
(c) for sharing creative practices, purposes and reaching common understanding when negotiating about embodied representations
(d) for the STEAM domains.

A plausible direction for the development of the participatory design methodologies seams to be a research through design (RtD) approach, an approach also known as design-based research [7] or action-based design [1]. Various researchers, e.g., [49], argue that the main reason that interaction design research has not always been successful is that it has not been guided by a sufficient understanding of the nature of design practice. Designers have to act on a complex design situation in all of its richness and complexity, and in a way that is appropriate for the specifics of the situation. This requires a designer to be fully immersed in the context of the case and to make sense of that context based on an understanding of the particular situation, and then to create an appropriate approach for the specific design task at hand ([49]: p. 61). The RtD approach has been already used in the design of interactive embodied learning systems (e.g. [2, 3]), involving both teachers and learners as co-designers and evaluators.

4.4 The Connection Between the Designed Multimodal Artifacts and the Online Makerspace

How will the products of the interdisciplinary design team will trigger the participation of teachers and students in the online makerspace? This comes to the dissemination preparation phase, where the domain experts together with learning psychologists, prototyping experts and teachers will have to document the needed hardware kits, to develop video instructions for the construction of the embodied learning environments and to provide learning scenarios for exploring the new interface modalities by teachers and students in and out of the classroom. The video instructions should be presented in a self-paced learning platform, should follow principles of video-based learning and should promote improvisation, adaptation and design freedom when the interfaces will be recreated by the students. The learning scenarios should be able to address various educational settings, from typical school-based learning to problem/scenario based learning, scaffolding and flipped learning approaches. Additionally, the learning scenarios should cope with both collaborative and individual learning and also consider gender issues with technology. These learning scenarios will serve as practical distilled demonstrators of the efficacy and impact of the embodied learning environments and also ensure the successful integration of the embodied learning environments in both distance and blended learning approaches. All these supportive material is of paramount importance both for the successful deployment of the designed interfaces in the teaching/learning practice and for their diffusion in schools all over the world.

5 Conclusions

The ultimate goal of learning and teaching is the acquisition of adaptive rather than routine expertise or competence i.e. the ability to apply knowledge and skills flexibly and creatively in a variety of contexts and situations (as opposed to simply completing stereotype school tasks without understanding) [17]. STEAM learning has to be conceived as an active, constructive and self-regulatory process of sense-making, understanding, and problem solving within a community of learners. There is ample evidence that instruction that endorses such a view needs to promote, from the very beginning, learner autonomy. This will create the necessary affective climate for the development of intrinsic and autonomous incentives for and authentic interest in the STEAM domains.

WeMake is a framework that pays particular attention to (a) how learners can be involved in and take the responsibility of the construction of their own learning devices, rather than in the mere use of pre-given learning devices, and (b) how teachers can be stimulated to enhance learners' con-constructive activities (c) how effective embodied learning environments can be disseminated in schools sooner than later (d) how to create a community of sharing resources about embodied learning environment with the participation of researchers, teachers and students. Essentially, WeMake proposes a work process for enabling students and teachers to build, exploit, share and assess their own embodied learning environments.

References

1. Abrahamson, D.: Seeing chance: perceptual reasoning as an epistemic resource for grounding compound event spaces. ZDM – Int. J. Math. Educ. **44**, 869–881 (2012)
2. Abrahamson, D.: Building educational activities for understanding: an elaboration on the embodied-design framework and its epistemic grounds. Int. J. Child-Comput. Interact. **2**(1), 1–16 (2014)
3. Antle, A.N.: The CTI framework: informing the design of tangible systems for children. In: Proceedings of the 1st International Conference on Tangible and Embedded Interaction, pp. 195–202. ACM (2007)
4. Antle, A.N., Wise, A.F.: Getting down to details: using theories of cognition and learning to inform tangible user interface design. Interact. Comput. **25**(1), 1–20 (2013)
5. Bakker, S., Antle, A., van den Hoven, E.: Embodied metaphors in tangible interaction design. Pers. Ubiquit. Comput. **16**, 433–449 (2011)
6. Bakker, S., van den Hoven, E., Antle, A.N.: MoSo tangibles: evaluating embodied learning. In: Proceedings of the Fifth International Conference on Tangible, Embedded, and Embodied Interaction, pp. 85–92. ACM (2011)
7. Barab, S., Squire, K.: Design-based research: putting a stake in the ground. J. Learn. Sci. **13**(1), 1–14 (2004)
8. Barfield, W., Caudell, T.: Basic concepts in wearable computers and augmented reality. In: Fundamentals of Wearable Computers and Augmented Reality, pp. 3–26 (2001)
9. Barsalou, L.W.: Grounded cognition. Ann. Rev. Psychol. **59**, 617–645 (2008)
10. Bevan, B., Gutwill, J.P., Petrich, M., Wilkinson, K.: Learning through STEM-rich tinkering: findings from a jointly negotiated research project taken up in practice. Sci. Educ. **99**(1), 98–120 (2015)
11. Bower, M., Sturman, D.: What are the educational affordances of wearable technologies? Comput. Educ. **88**, 343–353 (2015)
12. Buechley, L., Hill, B.M.: LilyPad in the wild: how hardware's long tail is supporting new engineering and design communities. In: Proceedings of the 8th ACM Conference on Designing Interactive Systems, pp. 199–207. ACM, August 2010
13. Carbonneau, K.J., Marley, S.C., Selig, J.P.: A meta-analysis of the efficacy of teaching mathematics with concrete manipulatives. J. Educ. Psychol. **105**(2), 380 (2013)
14. Coffman, T., Klinger, M.B.: Google glass: using wearable technologies to enhance teaching and learning. In: Proceedings of Society for Information Technology & Teacher Education International Conference, pp. 1777–1780, March 2015
15. Colella, V.: Participatory simulations: building collaborative understanding through immersive dynamic modeling. J. Learn. Sci. **9**(4), 471–500 (2000)

16. Cramer, E.S., Antle, A.N.: Button matrix: how tangible interfaces can structure physical experiences for learning. In: Proceedings of the Ninth International Conference on Tangible, Embedded, and Embodied Interaction, pp. 301–304. ACM, January 2015
17. De Corte, E.: Constructive, self-regulated, situated, and collaborative learning: an approach for the acquisition of adaptive competence. J. Educ., 33–47 (2011)
18. de Freitas, S., Levene, M.: Evaluating the development of wearable devices, personal data assistants and the use of other mobile devices in further and higher education institutions. In: JISC Technology and Standards Watch Report, (TSW030), pp. 1–21 (2003)
19. Dede, C.: Immersive interfaces for engagement and learning. Science **323**(5910), 66–69 (2009)
20. Dougherty, D.: The maker movement. Innovations **7**(3), 11–14 (2012)
21. Enyedy, N., Danish, J.A., Delacruz, G., Kumar, M.: Learning physics through play in an augmented reality environment. Int. J. Comput.-Support. Collaborative Learn. **7**(3), 347–378 (2012)
22. Goodman, S.G., Seymour, T.L., Anderson, B.R.: Achieving the performance benefits of hands-on experience when using digital devices: a representational approach. Comput. Hum. Behav. **59**, 58–66 (2016)
23. Gutwill, J.P., Allen, S.: Facilitating family group inquiry at science museum exhibits. Sci. Educ. **94**(4), 710–742 (2010)
24. Halverson, E.R., Sheridan, K.: The maker movement in education. Harvard Educ. Rev. **84** (4), 495–504 (2014)
25. Hatton, S., Birchfield, D., Megowan-Romanowicz, M.C.: Learning metaphor through mixed-reality game design and game play. In: Proceedings of the 2008 ACM SIGGRAPH Symposium on Video Games, pp. 67–74. ACM (2008)
26. Ishii, H., Ullmer, B.: Tangible bits: towards seamless interfaces between people, bits and atoms. In: Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems, pp. 234–241. ACM, March 1997
27. Kazemitabaar, M., Norooz, L., Guha, M.L., Froehlich, J.E.: MakerShoe: towards a wearable E-textile construction kit to support creativity, playful making, and self-expression. In: Proceedings of the 14th International Conference on Interaction Design and Children, pp. 449–452. ACM, June 2015
28. Kuznetsov, S., Paulos, E.: Rise of the expert amateur: DIY projects, communities, and cultures. In: Proceedings of the 6th Nordic Conference on Human-Computer Interaction: Extending Boundaries, pp. 295–304. ACM, October 2010
29. Lee, S.A., Chung, A.M., Cira, N., Riedel-Kruse, I.H.: Tangible interactive microbiology for informal science education. In: Proceedings of the Ninth International Conference on Tangible, Embedded, and Embodied Interaction, pp. 273–280. ACM, January 2015
30. Lindgren, R., Johnson-Glenberg, M.: Emboldened by embodiment six precepts for research on embodied learning and mixed reality. Educ. Researcher **42**(8), 445–452 (2013)
31. Lindgren, R., Moshell, J.M.: Supporting children’s learning with body-based metaphors in a mixed reality environment. In: Proceedings of the 10th International Conference on Interaction Design and Children, pp. 177–180. ACM, June 2011
32. Lindgren, R., Tscholl, M., Wang, S., Johnson, E.: Enhancing learning and engagement through embodied interaction within a mixed reality simulation. Comput. Educ. **95**, 174–187 (2016)
33. Link, T., Moeller, K., Huber, S., Fischer, U., Nuerk, H.-C.: Walk the number line – an embodied training of numerical concepts. Trends Neurosci. Educ. **2**(2), 74–84 (2013)
34. Lyons, L., Silva, B.L., Moher, T., Pazmino, P.J., Slattery, B.: Feel the burn: exploring design parameters for effortful interaction for educational games. In: Proceedings of the 12th International Conference on Interaction Design and Children, pp. 400–403. ACM, June 2013

35. Malinverni, L., Pares, N.: Learning of abstract concepts through full-body interaction: a systematic review. Educ. Technol. Soc. **17**(4), 100–116 (2014). doi:10.2307/jeductechsoci.17.4.100
36. Manches, A., O'Malley, C., Benford, S.: The role of physical representations in solving number problems: a comparison of young children's use of physical and virtual materials. Comput. Educ. **54**(3), 622–640 (2010)
37. Mpiladeri, M., Palaigeorgiou, G., Lemonidis, C.: FRACTANGI: a tangible learning environment for learning about fractions with an interactive number line. In: Proceedings of 13th International Conference on Cognition and Exploratory Learning in Digital Age (CELDA 2016), pp. 157–164 (2016)
38. Norooz, L., Mauriello, M. L., Jorgensen, A., McNally, B., Froehlich, J.E.: BodyVis: a new approach to body learning through wearable sensing and visualization. In: Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems, pp. 1025–1034. ACM, April 2015
39. Oehlberg, L., Willett, W., Mackay, W.E.: Patterns of physical design remixing in online maker communities. In: Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems, pp. 639–648. ACM, April 2015
40. Palaigeorgiou, G., Karakostas, A., Skenteridou, K.: FingerTrips: learning geography through tangible finger trips into 3D augmented maps. In: Proceedings of 2017 IEEE 17th International Conference on Advanced Learning Technologies (ICALT), pp. 170–172 (2017)
41. Peppler, K.A.: STEAM-powered computing education: using E-textiles to integrate the arts and STEM. IEEE Comput. **46**(9), 38–43 (2013)
42. Petrich, M., Wilkinson, K., Bevan, B.: It looks like fun, but are they learning? In: Honey, M., Kanter, D. (eds.) Design, Make, Play: Growing the Next Generation of STEM Innovators, pp. 50–70. Routledge, New York (2013)
43. Pouw, W.T., Van Gog, T., Paas, F.: An embedded and embodied cognition review of instructional manipulatives. Educ. Psychol. Rev. **26**(1), 51–72 (2014)
44. Price, S., Jewitt, C.: A multimodal approach to examining 'embodiment' in tangible learning environments. In: Proceedings of the 7th International Conference on Tangible, Embedded and Embodied Interaction, pp. 43–50. ACM, February 2013
45. Resnick, M., Martin, F., Berg, R., Borovoy, R., Colella, V., Kramer, K., Silverman, B.: Digital manipulatives: new toys to think with. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, pp. 281–287. ACM Press/Addison-Wesley Publishing Co., January 1998
46. Sarama, J., Clements, D.H.: "Concrete" computer manipulatives in mathematics education. Child. Dev. Perspect. **3**(3), 145–150 (2009)
47. Sherin, M., Jacobs, V., & Philipp, R. (Eds.). (2011). Mathematics teacher noticing: Seeing through teachers' eyes. Routledge
48. Sloutsky, V.M., Kaminski, J.A., Heckler, A.F.: The advantage of simple symbols for learning and transfer. Psychon. Bull. Rev. **12**(3), 508–513 (2005)
49. Stolterman, E.: The nature of design practice and implications for interaction design research. Int. J. Des. **2**(1) (2008)
50. Stull, A.T., Barrett, T., Hegarty, M.: Usability of concrete and virtual models in chemistry instruction. Comput. Hum. Behav. **29**(6), 2546–2556 (2013)
51. Taylor, K., Hall, R.: Counter-mapping the neighborhood on bicycles: mobilizing youth to reimagine the city. Technol. Knowl. Learn. **18**(1–2), 56–93 (2013)
52. Tolentino, L., Birchfield, D., Megowan-Romanowicz, C., Johnson-Glenberg, M.C., Kelliher, A., Martinez, C.: Teaching and learning in the mixed-reality science classroom. J. Sci. Educ. Technol. **18**(6), 501–517 (2009)

53. Vossoughi, S., Escudé, M., Kong, F., Hooper, P.: Tinkering, learning & equity in the after-school setting. In: Annual FabLearn Conference. Stanford University, Palo Alto (2013)
54. Wu, T., Dameff, C., Tully, J.: Integrating Google glass into simulation-based training: experiences and future directions. J. Biomed. Graph. Comput. **4**(2), 49 (2014)
55. Zacharia, Z.C., Olympiou, G.: Physical versus virtual manipulative experimentation in physics learning. Learn. Instr. **21**(3), 317–331 (2011)
56. Zimmerman, J., Forlizzi, J., Evenson, S.: Research through design as a method for interaction design research in HCI. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, pp. 493–502. ACM, April 2007

Onboarding Communities to the IoT

Mara Balestrini(✉), Gui Seiz, Lucas L. Peña, and Guillem Camprodon

Fab Lab Barcelona, Barcelona, Spain
{mara, gui, lucas, guillem}@fablabbcn.org

Abstract. With the advent of the Internet of Things, low-cost sensing technologies are becoming increasingly available, allowing citizens to collectively monitor and share data about the environment. A subset of these technologies is being made at maker spaces, using open source and affordable technologies. While these systems have the potential to power bottom-up participatory data networks, a key concern is that laypeople often fail to effectively setup and connect these sensors to the Internet because they lack technical skills and/or the systems' user experience is poorly designed. We present a novel onboarding application that aims to facilitate the process of sensor setup and connection by non-experts. It works by providing an integrated design experience, scaffolding the complexity of the process, and guiding the user in a conversational fashion. We hope to inspire other developers and designers to consider the needs of non-technical and diverse communities in the design of IoT systems.

Keywords: Digital social innovation platforms · Crowdsensing · Internet of Things · UX · Design

1 Introduction

In the past decade, the proliferation of open source technology such as Arduino, the creation of makerspaces and Fab Labs [13], and the growing popularity of crowdfunding platforms have fostered the design of new sensing technologies that allow citizens to collect and share all sorts of data, from temperature and humidity [3], to air quality or noise pollution [9] and nuclear radiation [14]. These new technologies aim to empower self-organised groups of citizens by providing Do-it-yourself (DIY) or more open systems that they can appropriate to tackle matters of concern, learn about technology and/or the environment. Notable examples are the Air Quality Egg (airqualityegg.com), Safecast (safecast.org) or RadiationWatch (radiation-watch.org).

While these novel and affordable technologies could foster the emergence of new forms of environmental monitoring, empowering citizens to collect, share and use data in their own terms, previous studies have revealed a number of technical and social issues that can hinder the appropriation of sensing practices at the grassroots level [8]. A key concern is that laypeople often fail to effectively setup and connect these sensors to the Internet because they lack technical skills and/or the systems' user experience is poorly designed [2]. Naturally, failing to set up and connect a sensor often leads to frustration and disengagement; and to a missed opportunity to contribute valuable data.

I. Kompatsiaris et al. (Eds.): INSCI 2017, LNCS 10673, pp. 19–27, 2017.
https://doi.org/10.1007/978-3-319-70284-1_2

How can low cost sensors be useful and engaging for laypeople? How can this technology, which was traditionally designed to be used by experts, be *enchanting* [15] for the general public? In this short paper we present a novel onboarding application for the Smart Citizen Kit, a crowdfunded open source sensor kit for environmental monitoring that was designed at Fab Lab Barcelona [9]. The application facilitates the process of sensor setup and connection by non-experts. It works by providing an integrated design experience, scaffolding the complexity of the setup process, and guiding the user in a conversational fashion.

We followed a participatory research through design approach to collaborate with community participants who were involved in a citizen sensing intervention aimed to tackle noise pollution in Barcelona. We here present how the process led to a novel open source application that can be readily appropriated by developers and designers working towards the design of engaging and user friendly IoT systems.

2 Background

Researchers and practitioners have studied many aspects of participatory sensing systems, from technology design and features to people's motivations to participate in citizen science and issues around data sensemaking [1, 5, 6, 24]. A large proportion of these findings refer to either citizen science projects or research studies where participants have been recruited to collaborate with experts and or facilitators. In consequence, there is little reference to the factors associated to the appropriation of participatory IoT and sensing tools at the grassroots level. How should such tools be designed for laypeople and grassroots communities to be able to autonomously use and profit from them?

It is important to consider that there are key differences between traditional Information and Communication Technologies (ICT) and the more novel sensing and IoT technologies. While personal computers and mobile phones are pervasive in everyday life, sensing technologies are still novel and largely unfamiliar to most people [8]. Until recently, sensing technologies tended to be embedded into existing products and the environment, which meant that the public had little access to them. This unfamiliarity with the technology and lack of skills to operate them can have an impact on how effectively people engage in data collection processes. Moreover, low-cost tools are sometimes still unreliable and hard to use [23].

Following participatory approaches to design novel sensing technologies can foster acceptance and appropriation. For example, DiSalvo et al. [8] introduced sensor and robotic technologies to residents in a neighbourhood, organising a set of activities such as scenario writing and mock-up development to inspire people to envision novel applications of the devices. As a result, the participants gained familiarity with the technology and appropriated it in ways that had been unanticipated by the researchers.

Next, we describe how we designed a new IoT onboarding application in collaboration with users. The application comprises assembling a Smart Citizen Kit device and its enclosure, configuring and naming it, and subsequently connecting it to the Internet to relay data to an online platform. The aim was to elicit people's preferences to co-create a system that makes it easy for users to engage with crowdsensing campaigns by learning how to deploy a low cost sensor. Furthermore, rather than trying to design a complete *plug&play* solution, we aimed to support the development of technical skills in users, by requesting that they assemble their own sensor and make sense of the setup process [19, 21].

3 Smart Citizen

Smart Citizen is an open source sensing platform that comprises a sensor kit (SCK), an online platform and a mobile application. The project was launched in 2012, instigated by the Fab Lab Barcelona [9]. The SCK consists of an Arduino-based electronic board and shield, a battery, a Wi-Fi antenna, a MicroSD card, and a set of sensors to monitor humidity, temperature, nitrogen dioxide, carbon monoxide, sound, and light. The SCK firmware is created to read, process and post this data, as well as battery charge level, intermittently upon setup. The online platform (smartcitizen.me) allows users to upload data from their SCKs, this is either automated through the kit or through CSV to share them online. Both the sensor kit and the online platform were developed with financial support from users through two crowdfunding campaigns.

From the 1200 SCKs that have been delivered to users around the world only a small percentage of these sensors (less than 20%) are kept online, contributing data to the Smart Citizen platform. Previous research has shown that users struggle to connect their SCKs to the Internet [2]. Informants explained that they found the process to be too long, confusing, and tedious. They also highlighted the need to have access to accessible documentation and troubleshooting advice to assist them while setting up their devices. While SCK does provide a place for such documentation, this experience is segmented and often complex for a non-technical audience (Fig. 1).

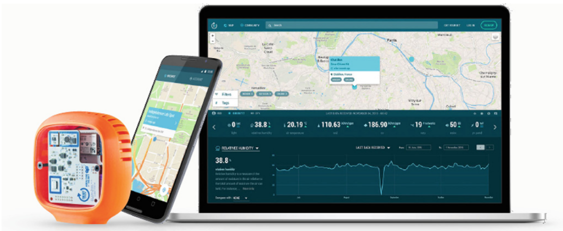

Fig. 1. The Smart Citizen kit, mobile app and data visualization platform.

4 Collecting User Requirements

We followed a participatory research through design approach to the development of a new onboarding application [25]. This method involves following consecutive iterations and evaluations to frame the problem and improve some characteristic of the studied phenomenon [12, 17], in our case, we focused on the onboarding experience to the Smart Citizen Kit by non-technical users. The aim was to collaborate with users to gather usability insights and co-design a new experience that was engaging and user-friendly, and allowed users to assemble a device, set it up and connect it to the Internet via Wi-Fi by themselves.

Three think aloud sessions with 27 community participants who were participating in a citizen science project but had never used a Smart Citizen Kit before were organised. This method requires that the participants talk aloud while performing a task [10]. Two sessions were individual (one participant in each), and a third was a group session with 25 participants who worked in groups of three. The sessions were held at the lab where participants were given pouches comprising a disassembled SCK, a USB cable and a 3D printed enclosure, and asked to assemble and install the sensors using their computers. Two observers took notes of how users performed the task and annotated verbatim. The collected data was analysed by two coders [4] and used as input during co-design workshops. The emergent issues were later synthesised under three themes: complexity, engagement, and hardware.

Complexity

Fragmentation: We identified that, if completed with no deviations, the installation process of the Smart Citizen Kit had a minimum of 20 tasks, some consisting of multiple steps. Often these tasks led to dead ends, or were largely unsuccessful for different reasons - e.g. lack of connectivity, no existing user account, etc., - which meant the user would have to restart the whole process.

Lack of signposting: Participants found the setup process to be *"too long"* and lacking status updates that helped them make sense of the stage of the process in which they were. Steps of the installation were seemingly out of context, asking users for potentially invasive information without an accompanying justification.

Engagement

We found that personal details such as exact location, name, email, etc. was the kind of information that participants least liked giving out and would most often question. Moreover, documentation was both difficult to find and difficult to troubleshoot. This often created confusion, disengagement and led to abandonment.

Language: Participants found the language used in the Smart Citizen platform too technical. They often found themselves wondering what key terms meant and confused on how to proceed with the installation process.

Hardware

The installation process required the use of a micro USB cable, which provided an obstacle for those who did not own one. This also made it difficult for people to install the device on a mobile device such as a phone or tablet, therefore limiting installation to a physical experience of using a computer.

5 A New Onboarding Experience

Two co-design workshops with two developers, one user experience designer and community participants were organised at the lab. During the workshops the themes resulting from the think aloud sessions were analysed and discussed by participants, designers and coders. The ambition was to design an onboarding experience that would increase the ratio of successful installations by reducing the complexity in signing up and fostering user engagement.

As a result, a number of prototypes were developed, tested and iterated in collaboration with community participants. We addressed complexity challenges by restructuring and condensing the sign up experience, harnessing the completeness effect [22] and introducing a progress bar. We fostered a sense of engagement by creating a sense of device ownership and developing a friendly visual and written language. Finally, we tackled hardware issues by creating a seamless experience where users do not need to use a USB cable or external assets during the installation process. Next, we present the onboarding application[1] and describe how we addressed these themes (Fig. 2).

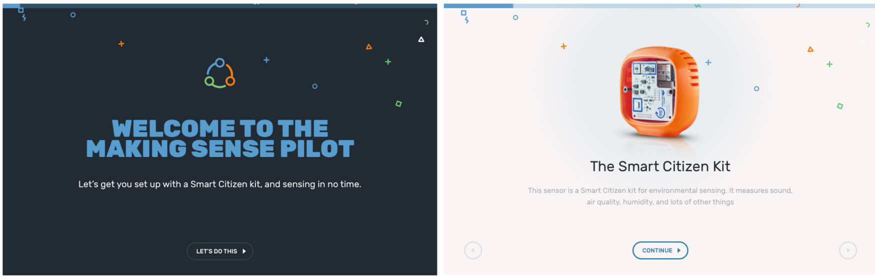

Fig. 2. Screen captures of the onboarding application.

5.1 Complexity

Restructuring and condensing the sign up experience. The installation sought to create context upfront and scaffold the complexity, leaving the most monotonous tasks towards the end. This took into account the Goal Gradient effect [7], whereby users efforts towards a goal increase as the proximity to the goal increases. The geolocalisation of the sensor, creation of account and submission of personal details were all moved to the end of the user journey, thus rendering the risk of a wiping out the effort if details were denied. The final, user detail, submission process essentially finishes the installation and redirects users to the Smart Citizen platform.

Completeness. The tasks were also compartmentalised in complexity. Each task was broken down into an appropriate number of steps that balanced ease of use with complexity, and difficulty with stagnation to create a continuous sense of flow [20] and progression. As each of the larger goals were broken down into steps, the larger goal was found to be less daunting, and easier to accomplish. By packaging these tasks as part of the larger goal we seek to provide a perspective of importance of the entire

[1] onboarding.making-sense.eu and https://github.com/MakingSense-EU/onboarding-app.

onboarding process. This harnesses the completeness effect [22], where by showing parts as being incomplete of a greater whole, more consumption of such parts takes place.

Progress bar. We paid special attention to visibility of system status and signposting. A progress bar, as well as constant feedback brought context to the tasks being undertaken and effectively broke the end goal down into achievable steps. The idea was to visualise each step (goal) as simply as possible to increase chance of completion [16].

5.2 Engagement

Creating a sense of device ownership. To foster a sense of ownership [2], and support the development of technical skills in participants, we designed the task of physically assembling the SCK device as part of the onboarding experience. This task did not exist in the previous onboarding system although SCKs are bought disassembled and instructions are hard to find. In the new onboarding experience, the assembly was outlined in the installation instructions and treated as importantly as the digital setup itself. Coupled with the ability to do seemingly trivial things such as naming the sensor, this assembly and ‘life giving’ to the device by the participants can contribute to fostering a sense of value, ownership and attachment to the physical device [18].

Creating a friendly language. Another improvement was in the written and visual language of the user journey. Previously the installation had been confusing due to the overtly technical nature of the language used, as well as the poor signposting. When components were referenced by name, participants had little choice but to guess as to which components were referred to. Simplifying the language, provided we educated when needed, can breed trust within the experience [11] and strengthen engagement.

The improved onboarding features a photographic and iconographic-based approach where each component and method of assembly is represented and needs to be accounted for before progressing. On each of the components a link for a more detailed description is provided-with the idea of familiarising participants with the basic technical components to facilitate the assembly and troubleshooting processes (Fig. 3).

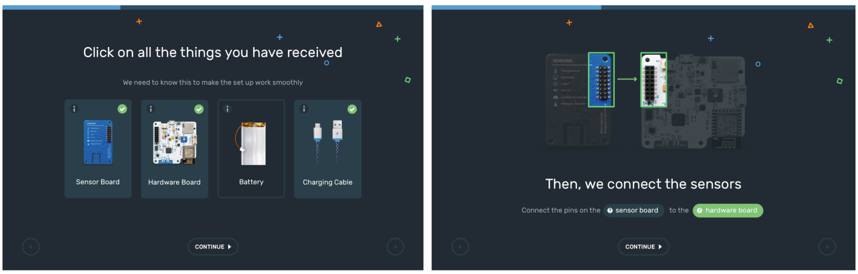

Fig. 3. Sensor hardware and assembly.

5.3 Seamlessness

Optimised for devices & removing the need for hardware. The process of syncing the SCK device with the platform was the hardest one to solve, as all the routes investigated had some degree of complexity or potential risk of failure. The original installation was executed via a USB cable, and the installation of a browser extension, via which the data was transmitted. With the aim of condensing the installation into a single process and reducing the reliance on collateral hardware, a mechanic was designed to transmit the device credentials to the platform using the embedded light sensor in the SCK sensor board.

The light sensor captures RGB (additive light values). Using this input we were able to divide the usable alphabet, numbers and symbols into grey-scale values. The light sensor is used as a way to easily measure grayscale. This makes it possible to transmit values representing characters, which make up the WiFi SSID and password the user previously enters to connect the sensor to their network. The handshake process itself begins when the sensor is waiting for configuration and is presented with the initialisation process. The sensor needs to be placed in front of the screen with the source of the grayscale handshake SVG area.

This approach allows for the potential of any device with a modern screen both matte and gloss to handshake with the sensor given the screen itself is placed directly against the light sensor. It is important to note that in order to minimise errors in this process, proximity and brightness of the screen are imperative to lower the risk of passing corrupted characters. The onboarding process was optimised for a variety of devices. Users can therefore benefit from the same experience whether they are on a desktop computer, a laptop or a touch device by using this seamless ‘light handshake’.

5.4 Tech Stack

When planning and building the technology stack, scalability through abstraction and portability was the goal. It was decided that the application should run in AngularJS providing itself as a JavaScript bundle upon page load. Unlike the previous onboarding this one is constructed separated from the SCK platform. This application serves to provide modularity through service oriented architecture (SOA) in application and appropriation through open source technology. Refactoring and releasing new developments on the onboarding can be scaled with this approach. A modular onboarding is entirely dependent on available and public APIs of its parent project, which drives developers of said project to open and document their endpoints.

Additionally, the application is structured as a wizard. There is one large HTML frame, providing heading and footer content accordingly, where the internal body structure is updated to present different steps required by the user. These steps, through an animation framework built for the project, ease and animate so as to give the appearance and experience of interactive page loads while simply calling JavaScript to update HTML. These body views are constructed in reusable code, where possible. These views are populated before page presentation, from a separate JSON array, and animated with respective easing, accordingly. This approach of code abstraction into a JSON object allows for individual manipulation of the page contents or views without the need to alter the other (Fig. 4).

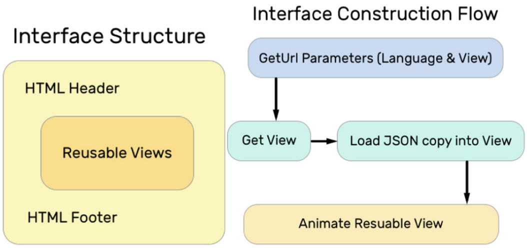

Fig. 4. Data structure.

6 Conclusion

New low-cost IoT devices are becoming more pervasive in everyday life, democratising access to data collection, sharing and analysis. With these tools people can join citizen science and sensing initiatives to address environmental threats and/or urban challenges. However, IoT devices are still largely unfamiliar for laypeople with low technical skills or often poorly designed in terms of usability. There is a need to make sense of user needs and requirements to design more usable systems that people can independently appropriate and use for their own purposes.

In this short paper we have presented a new onboarding application that aids the setup and installation process of a Smart Citizen Kit. Its design rationale aimed to address issues of complexity, engagement and hardware that were evidenced through think aloud usability sessions conducted with community participants who are part of a citizen science initiative. We found that following a participatory approach to design this application with and for communities was a fruitful method to elicit real needs and aspirations, and to create an engaging onboarding experience. The application is modular and open source to encourage external appropriation hoping to inspire other designers and developers of IoT systems.

Acknowledgement. We thank all our community participants and paper reviewers for their valuable contributions. The Making Sense project is made possible through EU H2020 CAPS ICT10a.

References

1. Bales, E., Nikzad, N., Quick, N., Ziftci, C., Patrick, K., Griswold, W.: Citisense: Mobile air quality sensing for individuals and communities Design and deployment of the Citisense mobile airquality system. In: PervasiveHealth 2012, pp. 155–158. IEEE, May 2012
2. Balestrini, M., Diez, T., Marshall, P., Gluhak, A., Rogers, Y.: IoT community technologies: leaving users to their own devices or orchestration of engagement? EAI Endorsed Trans. Internet Things **15**, 1 (2015)
3. Balestrini, M., Rogers, Y., Hassan, C., Creus, J., King, M., Marshall, P.: A city in common: a framework to orchestrate large-scale citizen engagement around urban issues. In: Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems, pp. 2282–2294. ACM, May 2017

4. Braun, V., Clarke, V.: Using thematic analysis in psychology. Qual. Res. Psychol. **3**(2), 77–101 (2006)
5. Canny, J., Mascolo, C., Paulos, E.: Pervasive analytics and citizen science. Pervasive Anal. Citizen Sci. **13**(2), 18–19 (2014)
6. Corburn, J.: Street Science: Community Knowledge and Environmental Health Justice. The MIT Press, Cambridge (2005)
7. Cryder, C.E., Loewenstein, G., Seltman, H.: Goal gradient in helping behavior. J. Exp. Soc. Psychol. **49**(6), 1078–1083 (2013)
8. DiSalvo, C., Louw, M., Coupland, J., Steiner, M.: Local issues, local uses: tools for robotics and sensing in community contexts. In: Proceedings of the Seventh ACM Conference on Creativity and Cognition, pp. 245–254. ACM, October 2009
9. Diez, T., Posada, A.: The fab and the smart city: the use of machines and technology for the city production by its citizens. In: Proceedings of TEI 2013, pp. 447–454. ACM, February 2013
10. Ericsson, K.A., Simon, H.A.: Protocol Analysis. MIT press, Cambridge (1993)
11. Franke, N., Piller, F.: Value creation by toolkits for user innovation and design: the case of the watch market. J. Prod. Innov. Manag. **21**(6), 401–415 (2004)
12. Gaver, W.: What should we expect from research through design? In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, pp. 937–946. ACM, May 2012
13. Gershenfeld, N.: Fab: the coming revolution on your desktop -from personal computers to personal fabrication. Basic Books (2008)
14. Kera, D., Rod, J., Peterova, R.: Post-apocalyptic citizenship and humanitarian hardware. Nuclear Disaster at Fukushima Daiichi: Social, Political and Environmental Issues, 97 (2013)
15. McCarthy, J., Wright, P., Wallace, J., Dearden, A.: The experience of enchantment in human–computer interaction. Pers. Ubiquit. Comput. **10**(6), 369–378 (2006)
16. Myers, B.A.: The importance of percent-done progress indicators for computer-human interfaces. In: ACM SIGCHI Bulletin, vol. 16, no. 4, pp. 11–17. ACM, April 1985
17. Nelson, H.G., Stolterman, E.: The design way: Intentional change in an unpredictable world: Foundations and fundamentals of design competence. Educational Technology (2003)
18. Norton, M.I., Mochon, D., Ariely, D.: The 'IKEA effect': when labor leads to love (2011)
19. Papert, S., Harel, I.: Situating constructionism. Constructionism **36**(2), 1–11 (1991)
20. Pearce, Jon M., Howard, S.: Designing for flow in a complex activity. In: Masoodian, M., Jones, S., Rogers, B. (eds.) APCHI 2004. LNCS, vol. 3101, pp. 349–358. Springer, Heidelberg (2004). doi:10.1007/978-3-540-27795-8_35
21. Resnick, M.: Behavior construction kits. Commun. ACM **36**(7), 64–71 (1993)
22. Sevilla, J., Kahn, B.E.: The completeness heuristic: product shape completeness influences size perceptions, preference, and consumption. J. Mark. Res. **51**(1), 57–68 (2014)
23. Snyder, E.G., Watkins, T.H., Solomon, P.A., Thoma, E.D., Williams, R.W., Hagler, G.S., Shelow, D., Hindin, D.A., Kilaru, V.J., Preuss, P.W.: The changing paradigm of air pollution monitoring (2013)
24. Willett, W., Aoki, P., Kumar, N., Subramanian, S., Woodruff, A.: Common sense community: scaffolding mobile sensing and analysis for novice users. In: Floréen, P., Krüger, A., Spasojevic, M. (eds.) Pervasive 2010. LNCS, vol. 6030, pp. 301–318. Springer, Heidelberg (2010). doi:10.1007/978-3-642-12654-3_18
25. Zimmerman, J., Forlizzi, J.: Research through design in HCI. In: Olson, J.S., Kellogg, W.A. (eds.) Ways of Knowing in HCI, pp. 167–189. Springer, New York (2014). doi:10.1007/978-1-4939-0378-8_8

Citizen Science Is in the Air – Engagement Mechanisms from Technology-Mediated Citizen Science Projects Addressing Air Pollution

Gavin McCrory(✉), Carina Veeckman, and Laurence Claeys

imec - SMIT, Vrije Universiteit Brussel, 1050 Brussels, Belgium
gmccrory19@qub.ac.uk, Carina.Veeckman@imec.be

Abstract. Environmental data is collected at unprecedented scales and speeds, targeting diverse societal challenges, and through the inclusion of multiple stakeholders. Yet, an understanding of enabling technologies involved in the engagement of citizens appear largely outside of the realm of air pollution. Recently, different air pollution projects have been rolled out in Europe and abroad; a structured analysis, however, of the way citizens are involved in these type of projects does not yet exist. In contribution to the ongoing EU-Funded project hackAIR, this paper therefore explores this research gap on the topic of air pollution and citizen science through the following question: *Which engagement mechanisms can be identified in existing air pollution citizen science projects?* We combine multiple literature sources, employ a systematic case study analysis and conduct seven qualitative interviews with key experts to target citizen science projects related to air pollution. Several mechanisms emerged at the interface between air pollution, citizen participation and knowledge production. These include: (1) Scale, (2) User-involvement and co-creation, (3) Communication, and (4) User motivation and aspects of behaviour. Despite its growing reputation in digital innovation, a majority of the mapped projects do not explicitly engage in any co-creation process. Multiple project insights suggest the importance of non-academic stakeholders as agents for communication and engagement. Campaign-based gamification can prove successful in establishing urgency in local contexts. Common engagement barriers include issues in the data contribution, science communication, technical project limitations, scaling and the critical nature of distributed sensors. This preliminary research offers a fruitful approach in assessing and comparing initiatives, and can enrich our understanding of the contribution that air pollution technology can have in citizen science.

Keywords: Citizen science · Air pollution · Open technology · Community sensing · Grassroots sensing

1 Introduction

Individuals and organizations are participating in new and unprecedented ways in the conduct of environmental research. This ‘participatory turn’ in science comes with a number of new roles and motivations of citizens, generating interest from both a

I. Kompatsiaris et al. (Eds.): INSCI 2017, LNCS 10673, pp. 28–38, 2017.
https://doi.org/10.1007/978-3-319-70284-1_3

practical and scientific perspective. This represents the holy grail in ICT innovation: broadening notions of citizenship and understanding the motives for long-term involvement of citizens in practices that contribute towards sustainability (Nov et al., 2011).

The wealth of citizen science opportunities taking place across the globe and the new technologies enabling robust environmental data appear largely outside of the realm of air pollution (Geoghegan et al. 2016). That is not to say that local air sensing is absent. Rather; analytical research in this context of air pollution remains underexplored at this point.

1.1 Research Questions

This goal of the reported research in this article is to identify initial trends and explore the mechanisms that are beginning to define this field. These findings will then feed directly into the development of the tailored engagement strategy that is being developed for hackAIR. This will be fulfilled via the following research question: *Which engagement mechanisms can be identified in existing air pollution citizen science projects?* This paper therefore contributes, through preliminary analysis as part of the EU funded hackAIR project, to this discussion.

1.2 The hackAIR Project

The hackAIR project[1] is an EU funded project underneath the framework of the H2020 ICT-11 call of the European Commission. This specific call for proposals has been initiated since 2012, with the aim of introducing “Collective Awareness Platforms for Sustainability and Social Innovation” (CAPS). As quoted by the European Commission (2016) “Collective Awareness Platforms are expected to support environmentally aware, grassroots processes and practices to share knowledge, to achieve changes in lifestyle, production and consumption patterns, and to set up more participatory democratic processes”. Web platforms are used within CAPS projects for enabling the dynamics of collective awareness construction, whereby the platform is a socio-technical solution that can be composed of multiple ICT tools, such as websites, social networks, data visualization, etc. The vision of CAPS projects is to work towards digital social innovation, whereby the innovative proposed solutions benefit social issues in society, and also engage society in developing the solution (Arniani et al., 2014).

Within this context, the hackAIR project is developing and pilot testing an open platform that enables communities of citizens to consume, monitor and publish data around outdoor air pollution levels - with specific focus on fine particulate matter PM10 and PM2.5. The open platform will collect measurements from existing air quality stations and open data, as well as user-generated sky-depicting images collected through a mobile application. Users will also be able to collect measurements from low-cost, open hardware sensors and will receive personalized air quality information about their

[1] http://www.hackair.eu/.

everyday activities. The hackAIR platform will be tested in two pilot locations, being Germany with the support of the environmental organization BUND and in Norway with the Norwegian Institute for Air Research (NILU). The hackAIR project envisions to engage more than 8000 individuals in total, and for which a specific engagement and behavioural change campaign will be designed. The engagement and behavioural change campaign will identify suitable tactics for engaging environmentally aware and unaware citizens about air pollution, and will measure the significant change in the way people act, think and feel about the air quality in their local neighbourhood before and after being involved with hackAIR.

2 Citizen Engagement

Engagement can be defined as the "interactive and iterative processes of deliberation among citizens and between citizens and government officials with the purpose of contributing meaningfully to specific public policy decisions in a transparent and accountable way" (Barney 2006). In such a context, engagement draws parallels with notions of participation and participatory democracy, as a condition to consider in the pursuit of more transparent collaborative government-citizen relations. However, when extended towards participation and the generation of knowledge and local-decision making, Poliakoff and Webb (2007) offer a wider understanding of citizen engagement as one where communication occurs outside of the "ivory tower". This type refers to more recent scientific endeavours that move beyond those led by experts with specialist knowledge, and towards science that is more participatory in nature. For example, by integrating mechanisms to initiate engagement into spaces of academia and government, they touch upon broader processes of institutionalization that can foster citizen engagement in science and policy. Such mechanisms should balance short term and long-term objectives, blend face-to-face and virtual interactions, encourage deliberation and dialogue between different stakeholders, adhere to principles of transparency and trust, facilitate citizen involvement in project goals (Powell and Colin 2008). Given recent momentum towards approaches and paradigms that advocate for dispersed decision making (Certomà et al. 2015; Olsson et al., 2006), citizen science is increasing in its appeal and application.

2.1 Citizen Science and Environmental Monitoring

Kruger and Shannon (2000) define citizen science as "the process whereby citizens are involved in science as researchers". With an established tradition in Ornithology and Astronomy, and examples traced back to the early 20th century, participation between experts and non-experts might not necessarily be nascent (Bonney et al. 2009; Geoghegan et al. 2016). Recently however, citizen involvement in environmental monitoring has emerged as a broad, dynamic and ever-evolving set of practices whereby citizens can now engage with, contribute towards and direct multiple processes of scientific enquiry (Alfonso et al., 2015; Fritz et al., 2009; Higgins et al., 2016). Further, the prevalence of smartphones, decentralized sensor nodes and digital platforms now mean that citizen contribution to science is more dynamic than ever before.

Environmental data is being collected at unprecedented scales and speeds; towards the fulfilment of diverse societal challenges, and through the inclusion of multiple stakeholder sets. For example, the integration of multiple sensors into smartphones (GPS, Accelerometer etc.) has enabled remote or rural communities to partake in participatory sensing opportunities that are directly related to local concerns (Bonney et al., 2009).

Despite the benefits of citizen science, citizen participation remains in itself central to the success of a project on two levels. Firstly, initiatives are often constrained by the voluntary contributions of citizens, a majority of whom have different skill levels and expectations to contribute (Haklay 2015). Secondly, science projects remain limited by their ability to both involve and engage citizens over the extended course of a research project.

3 Methodology

3.1 Research Design

This research approach can be defined through three phases: (1) Systematic collection of citizen science projects on the topic of air pollution, (2) identification of projects for qualitative interview by narrowing down the initial list of projects, based on specific criteria in line with the goals of hackAIR, and (3) qualitative interviews with key actors of the selected projects. The research combines a literature study, systematic analysis of the projects, and seven qualitative interviews to target citizen science and air pollution. An analysis of the projects was adopted to accommodate an oversight of projects with a focus on air pollution in sufficient breadth. Subsequently, qualitative interviews were then conducted to allow for a deeper investigation of seven projects. This methodological triangulation guarantees a comprehensive study into the key determinants in this field (Yin 2015).

First, a systematic approach to collect projects was employed for this research. Table 1 displays the four search strings that were entered into Scopus database to retrieve a core empirical base.

Table 1. Systematic search strings for case-study analysis

Keywords	Results
"Citizen Science" AND "Air Quality"	N = 25
"Participatory" AND "Air"	N = 425
"Citizen Science" AND "Air Pollution"	N = 25
"Citizen Science" AND "Air"	N = 67

Upon completion, the abstract of each paper was reviewed with the following criteria: (1) *Does this paper touch upon citizen science and air pollution?* (2) *Does this paper have an explicit approach to citizen participation?* Additionally, Google scholar and a snowball approach from existing literature (both grey and peer-reviewed) were used to extend the list of projects for analysis. Overall, this process generated 35 projects for analysis. Secondly, additional criteria were applied to identify projects suitable for qualitative interviews. These were as follows: *Technology* – the project must involve technology for data collection, aggregation or communication; *Citizen engagement* – Citizens must be involved in at least one stage of the project; *ICT* – a web platform or companion app must be available. When applied, these criteria yielded six core projects for investigation: (1) EveryAware APIC, (2) Urban AirQ, (3) CITI-Sense, (4) iSPEX, (5) SecondNose and (6) ClairCity. With its focus on co-creation and local level air pollution, a supplementary interview was conducted with an expert from iMinds Living Labs with respect to the CityZen project. Specific details about interviews, respondents and projects can be seen in Annexes 1 and 2.

3.2 Qualitative Interviews and Data Analysis

Often under-documented in the scientific process, systematization and transparency in the data analysis process can help in reconciling limitations to qualitative research i.e. authenticity, credibility and reproducibility. In the context of this paper, skype interviews with seven key actors of mapped projects were conducted between April 18th and May 9th, 2017. Interviews followed a flexible tone and a fluid structure, focusing on broad themes rather than a more rigid approach. All interviews lasted between 45–60 min in length, were recorded with the consent of interviewees, and were fully transcribed.

Subsequently, transcriptions were sent to interviewees for confirmation in order to eliminate potential biases and misinterpretations that might have occurred. Comprehensive transcripts were then analyzed with the use of qualitative analysis software MAXQDA. Qualitative codes were derived using key parameters in the scientific field related to engagement, and were deliberately treated in relation to other sources of data, such as project deliverables, official reports and peer-reviewed research. Coded text varied relatively in length, but was never limited to individual words. Rather, phrases and sentences were used to effectively connect qualitative insights to the context in which they were generated. Finally, the analysis of findings from qualitative interviews was corroborated with the general trends uncovered in the systematic project mapping. The following section will delineate preliminary insights from this process.

4 Results

Based on the analysis of findings, interlinked with general trends of systematic project mapping, several mechanisms emerged at the interface between air pollution, project participation and knowledge production. In summary, common mechanisms include: (1) Scale, (2) User-involvement and co-creation, (3) Multiple facets of communication, and (4) Linkages between user motivation and aspects of behaviour.

4.1 Scale

Findings suggest that scale is a central determinant to the process of citizens' engagement. Interestingly, the introduction of ICT technologies is often heralded as a catalyst in the expansion of citizen science initiatives. Investigated projects utilize ICT in different ways, ranging from solely data access to more complex forms of data contribution. Projects such as iSPEX and EveryAware achieved unprecedented quantities of user-generated data during their period, whereas Urban AirQ deliberately focuses on an acute, local air pollution issue. Both approaches have discrete advantages and limitations, and as such, interact differently with aspects of scale. Scaling represents processes of replication and diffusion of innovation across spatial, temporal or institutional scales. iSPEX achieved critical mass in terms of data collection, but the institutional and technical capacities of the project created a challenging environment to replicate the project beyond iSPEX EU. As affirmed by Frans Snik, the lead researcher in the iSPEX project *"iSPEX reached a scope where we had to transition to a more professional, operational state than a university could support"*. The Urban AirQ project remained at a confined spatial scale, creating a managed environment for engagement and citizens. In fact, spatial scale was the largest determinant in Urban AirQ; *"We knew that topic of air quality was urgent to citizens [in Valkenburgerstraat], and we invited people from this street to come to the kick-off meeting"* – Gijs, Waag Society.

4.2 Co-creation

Co-creation is rapidly gaining traction as a broad concept that potentially confronts the often-exclusive nature of digital innovation. It represents a set of practices that facilitate the rapid collection and integration of qualitative insights into innovation processes. In terms of benefits, co-creation processes fuse multiple forms of knowledge, ranging from technical to environmental expertise. Furthermore, co-creation is underpinned by the assertion that all citizens have the right to, and are able to, contribute towards innovation processes (Sanders, 2002). A majority of the mapped projects (20/35) do not explicitly engage in any co-creation process. It is quite common that digital services and solutions are designed in contexts that exclude user participation until the beginning of the testing period. This results in a situation where the needs of users are not adequately captured during the formative stage of product or service development. In the context of a digital platform or social innovation, such exclusion can result in minimal user input regarding the user interface, usability the platform, look and feel of the sensor and user concerns. Such concerns are pivotal in the success of any innovation and ultimately reflect both user adoption and innovation diffusion potential.

Importantly, co-creation activities offer spaces for participation that are currently still lacking in the context of air pollution (Sîrbu et al. 2015). This trend is also emphasized by the broader assertion that, in citizen science research, user roles are typically consigned to data collection (Geoghegan et al. 2016). Although projects mapped in this project largely mirror this trend, there are some notable exceptions. SecondNose facilitated a co-creation process that directly fed into interface and application development. Due to funding issues however, this process was cut short before completion. Urban AirQ developed an engagement strategy that integrated co-creation at the stage of ideation and problem development. This ultimately led to a research design that was developed by, and owned by, non-experts.

Whilst not explicitly linked to co-creation but rather participation in science, CITI-SENSE encouraged the integration of subjective sensing, perceptions of air quality and alternative interpretations of air pollution. Each of these projects, in different ways, suggest that steps are being taken in the field of air pollution to move beyond traditional forms of knowledge and practice. This sentiment is encapsulated by Frans Snik, who states that "*I still firmly believe that the next big breakthrough in terms of citizen science, outside Astronomy and classification, will be with respect to air pollution*".

4.3 Communication

Findings highlight that communication interacts with citizens' engagement on multiple different levels. Firstly, ICT communications create digital avenues for projects to directly communicate with various stakeholders. For example, for Urban AirQ and CITI-SENSE, users and other citizens were directly able to access data measurements using web or companion applications. Second Nose and EveryAware APIC both used mobile sensors, which were dependent on a smartphone to send and communicate collected measurements. iSPEX stands out as a unique example in the field of air pollution to this point; the use and analysis of mobile images as proxies for outdoor air pollution.

Multiple project insights suggested the importance of non-academic stakeholders as agents for communication and engagement. This was evident in different city campaigns as part of the EveryAware project: *"It is important to have people that are experts in engagement working alongside engineers. We could see a clear difference in the number of citizens interested in Turin (without), and the amount of people in London (with). It is not easy to communicate with users, and we need to test the technology well without giving it away"* - Alina Sirbu, a computer scientist in EveryAware. In projects including EveryAware and Urban AirQ, separate non-state and non-academia actors were responsible in prominent aspects of user engagement.

Qualitative insights suggest that currently, academic institutions often assume roles that extend beyond traditional research. These operate at different levels, from the level of a university that can transfer knowledge towards other partners, to a researcher with expertise in science communication to engage directly with citizens and users. Indeed, projects indicated that users respond positively to a visible presence from academia, particularly for processes such as data quality, acknowledgement of contributions, translating input into impact, and offering support.

Projects differed in terms of their identification of target groups. Generally, however, and given the largely researcher-heavy orientation of citizen science, target groups such as universities and schools were considered strategic options for engagement. This is evident in projects such as AirBezen, EveryAware and CITI-SENSE. In the context of CITI-SENSE, children were identified for two reasons: (1) they represent a demographic that is commonly considered the most vulnerable to outdoor air pollution, partly due to their prolonged exposure outdoors. Secondly, they offer an avenue to resonate with parents, which can prove effective in fostering new ways of thinking about air pollution.

4.4 Motivations and Behaviour Change

As established in literature, there are multiple motivations of users engaging in citizen science (Rotman et al., 2014). Project mapping highlights that, in the context of air pollution and enabling technologies, research into the relationship between user motivations and behaviour remains inconclusive. Interview findings, complemented by academic contributions, do offer insights into user motivations and behaviors. Sources suggest that identified motivations for users of air pollution sensors are incredibly diverse, ranging from curiosity, interest in contributing to science and interest in local pollution, to using data to inform policy-making and encourage new forms of governance (Land-Zandstra et al., 2016; Leonardi et al., 2014).

For example, the EveryAware project demonstrates that, in the context of air pollution, campaign-based gamification proved successful in establishing urgency in local contexts. Interestingly, behavioural research conducted as part of the project actually found that, rather than leading to increased avoidance of pollution hotspots, users were drawn to high-pollution areas due to their motivation in contributing to science, and curiosity in the accuracy of the sensors (Sîrbu 2015). Drawing parallels with other strands of citizen science research, gamification is questionable as a long-term solution for all users in the context of pro-environmental behaviour (Geoghegan et al., 2016). Furthermore, lessons can be drawn from Leonardi et al. (2014), who identify fundamental differences between the motivations of volunteers for initial and long-term participation in the SecondNose project. This is evident in the shift away from curiosity about the accuracy of sensors and relevance to their own habits and everyday locations, and towards the need to cross-check 'subjective sensing' with more accurate or precise readings.

5 Conclusion

Through a multi-method approach, the aim of this research was to provide insights into the growing field of citizen science, with a particular focus on practices that aim to reduce the negative contributions of air pollution in Europe.

As expressed above, scale and scaling are fundamental issues that must be considered when establishing a vision on citizens' engagement. Given the appeal in attracting a large number of participants, the usage of ICT can be indispensable in generating collective forms of awareness. These tools and projects are however embedded within overarching institutional and organizational contexts, both of which can ultimately impact ICT diffusion. Co-creation with citizens is currently not always at the core of the analysed projects. Although co-creation as a concept is often used, the participation of citizens remains, in different projects, limited to data gathering only. In addition to this, classical forms of data are most commonly used. Despite these assertions, some signs are promising; alternative interpretations of AQ measurements are increasing and there are examples of projects that integrate subjective sensing into environmental monitoring. This preliminary research offers a fruitful approach in assessing and comparing initiatives, which can subsequently enrich our understanding of engaging citizens in scientific practice.

Appendix 1: Selected Interview Information

Project	Expert	Organization	Conducted
CityZen	Bas Baccarne	imec Living Labs	18/04/17 10:00
Urban AirQ	Gijs Boerwinkel	Waag Society	13/04/17 10:00
Clair City	Dr. Enda Hayes	UWE Bristol	21/04/17 14:00
CITI-SENSE	Sonja Grossberndt; Hai-Ying Liu	NILU	26/04/17 10:00
Second Nose	Chiara Leonardi	Fondazione Bruno Kessler	28/04/17 10:00
iSPEX	Dr. Frans Snik	University Leiden	28/04/17 11:00
EveryAware: APIC	Dr. Alina Sirbu	University of Pisa	08/05/17 11:00

Appendix 2: Qualitative Interview Topic List

Section A: General Project Information

1. Can you tell me a bit about the general aims and ambitions of your project?
2. More specifically, can you explain a bit about the role of yourself and your organization in the project?
3. Could you describe how community engagement fitted within your strategy?

Section B: Broad engagement

4. It would be great if you could tell me a bit about the process that your project took towards narrowing and defining target groups?
5. Likewise, could you describe how your engagement process was influenced by the different technologies involved in your project?
6. Could you tell me a bit about the roles and responsibilities of project partners in implementing the engagement strategy?

Section C: Characteristics of engagement

7. Could you explain a bit how your engagement strategy was tailored to reflect multiple stages of engagement?
8. Who was responsible for the communication of materials towards target groups for initial selection? This can be both via online/offline activities.
9. Furthermore, throughout your project were there any periods where corrective actions needed to be taken, or when you needed to adapt to changing circumstances?

Section D: Engagement mechanisms

10. Could you tell me a bit about how your project provided feedback to the local community of their contribution?
11. How were results relayed towards citizens in the local area? And by whom?
12. Did your project think it was important to provide recognition to citizens engaged in the project?
13. Was there a gamification or reward component to your project? How were incentives considered to keep citizens engaged and motivated?

Section E: Moving Forward

14. If you could pinpoint two main challenges that your project experienced that are directly related to engaging users for your project, what would they be?
15. What advice would you give to other projects that can help then in approaching an engagement strategy?

References

Alfonso, L., Chacón, J.C., Peña, G.: Allowing citizens to effortlessly become rainfall sensors. In: E-Proceedings of the 36th IAHR World Congress, 28 June – 3 July 2015, The Hague, The Netherlands, (1), pp. 1–5 (2015)

Arniani, M., Badii, A., Liddo, A., Georgi, S., Passani, A., Piccolo, L.S.G., Teli, M.: Collective awareness platforms for sustainability and social innovation: an introduction. (2014). http://booksprints-for-ict-research.eu/wp-content/uploads/2014/07/BS5-CAPS-FIN-003.pdf

Barney, D.: The morning after: citizen engagement in technological society. Techne **9**(3), 23–31 (2006)

Bonney, R., Cooper, C.B., Dickinson, J., Kelling, S., Phillips, T., Rosenberg, K.V., Shirk, J.: Citizen science: a developing tool for expanding science knowledge and scientific literacy. Bioscience **59**(11), 977–984 (2009)

Certomà, C., Corsini, F., Rizzi, F.: Crowdsourcing urban sustainability. Data, people and technologies in participatory governance. Futures **74**, 93–106 (2015). doi:10.1016/j.futures.2014.11.006

European Commission. CAPS Projects in FP7 (2016). https://ec.europa.eu/digital-single-market/en/programme-and-projects/caps-projects-fp7

Fritz, S., McCallum, I., Schill, C., Perger, C., Grillmayer, R., Achard, F., Kraxner, F., Obersteiner, M.: Geo-wiki.org: The use of crowdsourcing to improve global land cover. Remote Sens. **1**(3), 345–354 (2009) http://doi.org/10.3390/rs1030345

Geoghegan, H., Dyke, A., Pateman, R., West, S., Everett, G.: Understanding Motivations for Citizen Science, May 2016

Haklay, M.: Citizen Science and Policy: A European Perspective, 4 (2015)

Higgins, C.I., Williams, J., Leibovici, D.G., Simonis, I., Davis, M.J., Muldoon, C., van Genuchten, P., O'hare, G., Wiemann, S.: Citizen observatory web cobweb a generic infrastructure platform to facilitate the collection of citizen science data for environmental monitoring. Int. J. Spatial Data Infrastruct. Res. **11**, 20–48 (2016) http://doi.org/10.2902/1725-0463.2016.11.art3

Kruger, L.E., Shannon, M.A.: Getting to know ourselves and our places through participation in civic social assessment. Soc. Nat. Resour. **13**, 461–478 (2000)

Land-Zandstra, A.M., Devilee, J.L.A., Snik, F., Buurmeijer, F., van den Broek, J.M.: Citizen science on a smartphone: participants' motivations and learning. Pub. Underst. Sci. (Bristol, England), **25**(1), 45–60. http://doi.org/10.1177/0963662515602406

Leonardi, C., Cappellotto, A., Caraviello, M., Lepri, B., Antonelli, F.: SecondNose: an air quality mobile crowdsensing system. In: Proceedings of the 8th Nordic Conference on Human-Computer Interaction Fun, Fast, Foundational – NordiCHI 2014, pp. 1051–1054 (2014). http://doi.org/10.1145/2639189.2670273

Nov, O., Arazy, O., Anderson, D.: Technology-mediated citizen science participation: a motivational model. In: Proceedings of the Fifth International AAAI Conference on Weblogs and Social Media, pp. 249–256, July 2011 http://doi.org/10.1145/1940761.1940771

Olsson, P., Lebel, L., Garden, P., Berkes, F., Adger, W. N., Cash, D. W., Pritchard, L., Young, O.: Scale and cross-scale dynamics: governance and information in a multilevel world. Ecology and Society, **11**(2), 8 (2006). http://uu.summon.serialssolutions.com/2.0.0/link/0/eLvHCXMwtV1La9wwEBbbQKCX0jYpSR-gUy_GwSvZlhToYZPmQcmLPC65LNYrmKTe4Oz-_44kv3bJoQ3pxdjC2Ov9Ps9o5G9mEKJkJ4lXbIIEN6ETZrlOhVHEFQrTlNBc0hwcsq-4cXrLJ2f81yk_GI3ayhX92H8FHsYAepdI-w_gdxeFAdgHCsAWSADbv6LBFWAQvhDsO38Y

Poliakoff, E., Webb, T.L.: What factors predict scientists' intentions to participate in public engagement of science activities? Sci. Commun. **29**, 242–263 (2007)

Powell, M.C., Colin, M.: Meaningful citizen engagement in science and technology what would it really take? Sci. Commun. **30**(1), 126–136 (2008). http://doi.org/10.1177/1075547008320520

Rotman, D., Hammock, J., Preece, J., Hansen, D., Boston, C., Bowser, A., He, Y.: Motivations affecting initial and long-term participation in citizen science projects in three countries. iConference, 110–124 (2014). http://doi.org/10.9776/14054

Sanders, E.B.-N.: From user-centered to participatory design approaches. In: Design and the Social Sciences: Making Connections, pp. 1–8 (2002)

Sîrbu, A., Becker, M., Caminiti, S., De Baets, B., Elen, B., Francis, L., Gravino, P., Hotho, A., Ingarra, S., Loreto, V., Molino, A., Mueller, J., Peters, J., van den Bossche, J.: Participatory patterns in an international air quality monitoring initiative. PLoS ONE, **10**(8), 1–19 (2015). http://doi.org/10.1371/journal.pone.0136763

Yin, R.K.: Qualitative Research from Start to Finish. Guilford Publications, New York City (2015)

Community Based Initiatives and New Communication Technologies: A Preliminary Analysis Towards an Overall Assessment

Alessandra Prampolini(✉) and Antonella Passani

T6 Ecosystems S.r.l, Rome, Italy
a.prampolini@t-6.it

Abstract. Grassroots initiatives, initiated and managed by local communities, have been spreading in cities all around Europe, addressing the citizens' needs and the daily challenge of finding sustainable solutions for urban environments about food, waste, transport or energy. In recent years, these communities are increasingly making use of new technologies and some of them are currently shaped around Internet platforms, which became critical to their existence. This paper address them as grassroots online initiatives, aiming to outline their specificities and to discuss the methodology to research their impacts, with a specific attention towards the environmental dimension of their contribution to the society. The analysis is descriptive in nature and it has the main goal of integrating the emerging literature on this topic and of paving the way for further research activities in the field. To reach this objective, it presents a review of how new communication technologies have allowed the development of different kinds of community based initiatives, how these initiatives can be assessed and if they present significant differences with respect to the more traditional, face to face ones, focusing on a selection of case studies around Europe. An in-depth analysis of two initiatives, dealing respectively with sustainable mobility and sustainable food production, allows presenting some main findings.

Keywords: Community-based initiatives · Online platforms · Impact assessment · Low-carbon

1 Introduction

The need for a transition towards a low-carbon society is currently widely acknowledged and at the top of many national and international agendas [14]. Policy makers, businesses, administrations at different levels and civil society are all involved in the processes required to trigger and spread new solutions for sustainable lifestyles. Among these actors, in recent years the role of local communities and grassroots action has gained relevance around Europe, filling in the gap between local and national and international levels of governance. Solutions generated and developed at local level are progressively complementing higher-level strategies, providing governments the opportunity to work on new inputs and solutions based on community needs and available local resources [3, 10, 14]. Those solutions can be both social or

I. Kompatsiaris et al. (Eds.): INSCI 2017, LNCS 10673, pp. 39–51, 2017.
https://doi.org/10.1007/978-3-319-70284-1_4

technological in nature, and the most significant innovations are often generated at the intersection of the two perspectives [12]. From 2013 until 2016, the European research project TESS[1] mapped and analyzed dozens of community-based initiatives active in six countries around Europe (Germany, Italy, Spain, Finland, Scotland, Romania), dealing with four sector and the relative social challenges: food, waste, transport and energy. TESS defined community-based initiatives towards a low-carbon society as "collective actions which aim to mitigate greenhouse gas (GHG) emissions and reduce other adverse environmental impacts", where "a community is a group of people who (feel that they) share something in common, be that a connection through interest, place, lifestyle, culture or practice. They can include a combination of individuals, institutions and businesses, which instigate and deliver low-carbon innovations in partnership and in one or more domains (e.g. water, transport, energy, food)[2]."

This paper presents a preliminary investigation about how new communication technologies are triggering the development of different kinds of community based initiatives, how these initiatives can be assessed and if they present significant differences with respect to the more traditional, face-to-face ones. As the paper illustrates, the quick growth of online community initiatives makes their investigation relevant and necessary to understand the evolution of grassroots activation and impacts. Moreover, the specificities of new ICT tools are shaping the same idea of community and actions in new ways still under-investigated. The analysis represents an integration to the core topic of research of the project and for these reason it is based on a limited dataset, descriptive in nature, aiming to pave the way for future research activities looking at the interdependencies between online and offline community-based activities[3].

The paper is structured as follows: Sect. 2 defines the online community-based initiatives, describes their mapping and illustrates the selection process to identify the two interviewed initiatives; Sect. 3 outlines the assessment framework according to the online initiatives characteristics; and Sect. 4 presents two selected case studies, illustrating the main findings and highlighting open research questions.

2 Community-Based Initiatives and Online Community-Based Initiatives: Theoretical Background and Sampling

Community-based online initiatives can be broadly defined as online initiatives, social media campaign and other progressive networks that connect and synchronize the agendas of different actors involved in community-based projects and/or creating virtual online communities experimenting sustainable life-styles. They create novel channels for information and counter-information focused on sustainability and can organize and synchronize the initiatives of various social actors (e.g. sustainable

[1] Seventh Framework Programme, Grant Agreement No. 603705 (Project TESS).

[2] From TESS Description of Work.

[3] This paper is based on the results presented in TESS D4.2, "Multi-criteria analysis for carbon efficient projects": http://www.tess-transition.eu/wp-content/uploads/2016/11/TESS_D4.2_Multi-criteria-analysis-for-carbon-efficient-projects.pdf.

products and consumption behavior)[4]. This highly comprehensive definition can be narrowed according to the following criteria:

- online initiatives should have internet as their main "location": every grassroots initiative may have a website, but this would not make them an online initiative. An online initiative has born and grown around an online platform (website, social network, etc.) and could not act without it;
- online initiatives can have a wider geographical scope then non-online ones and they are often located in more than one region/country. This research considers initiatives that are based in Europe, in one or more of its countries/regions. A few examples originated in other parts of the world but they have a significant diffusion across Europe;
- in addition, online initiatives should match the following main criteria: they should be initiated and managed by communities (grassroots), they may have received public money but they should not be started up by public administrations, they may be non-for profit as well as for profit, but the revenue model should serve the community and they should belong to one of the four domains: food, transport, energy and waste.

About the "grassroots" nature of the initiatives, it is useful to underline that the concept of "community" changes when talking about online social interaction. The mainstream definition requires that a community has all or some of the following characteristics: territorial proximity, close and informal relationships, mutuality, common values and beliefs, organized interactions, strong group identity, cultural similarity. When online, communities may have common values and develop a shared identity but this is not always the case, and territorial proximity and close relationship may be substituted by frequent informal interaction based on common interests.

In this regard, a quite exhaustive definition may be the one formulated by Leimeister et al. [7]:

> "A virtual community consists of people who interact together socially on a technical platform. The community is built on a common interest, a common problem, or a common task of its members that is pursued on the basis of implicit and explicit codes of behavior. The technical platform enables and supports the community's interaction and helps to build trust and a common feeling among the members."

Platforms can attract participants in two ways: through a community-based approach or through a crowdsourcing one. In the first scenario the aggregation of the online community builds on a pre-existing, offline one, that makes the platform its main tool for action. Regarding the second scenario, there isn't a univocal definition of crowdsourcing, still it's possible to agree on some basic features. It is an online process that involves the participation of the crowd, gathered around a vision, a claim and/or a common target. In this case, participants do not previously know each other before joining the platform and they can keep acting together without ever having face-to-face contacts. Often it is a form of outsourcing of a task that would be not possible to handle internally at enterprise or research level. Moreover, the undertaking of a task by the

[4] From TESS Description of Work.

crowd entails mutual benefit, for the user and for the crowdsourcer. The users can benefit from various aspects: economic, social recognition, self-esteem and more [2].

The mapping of the representative selection of online initiatives around Europe operated on all the criteria illustrated above and brought to the identification of 15 initiatives. The small number of initiatives reported in the mapping depends on one significant factor: most of the online initiatives identified within a first, wider mapping, once further investigated revealed to be not grassroots in their nature, meaning that an institution or more often a private company started and still managed them. Despite the wide participation of local or international communities and their ethical goals, such initiatives have not been included in the final mapping, below:

Based on a desk research, these initiatives present the following characteristics:

- Domain: five are from the Food domain, two from the Transport domain, four from the Energy domain and four from the Waste domain.
- Goal: all the initiatives present a clear goal driven by ethical motivations, aiming to tackle social, economic and/or environmental challenges.
- Geographic scope: six initiatives are originating from a European country, which correspond also to their geographical boundary; seven initiatives, started in a European country, have today a European or international scope; while two have started out of Europe but were equally selected for their representativeness and diffusion also within the European context.
- Members: only for seven of the initiatives make available data about the number of their members. When available, the number can vary extremely, going from a few dozens to many millions.
- Among the 15 initiatives, 8 are community-based while 7 use a crowdsourcing approach.
- In the selection, for profit initiatives have always been excluded, and two top-down initiatives have been included because the organisation that started the initiative was grassroots and/or non-for-profit in itself (n. 4 – Roadsharing and n. 7 – Lifegate Impatto Zero).
- One initiative gives a profit to the shareholders (n. 7 – Lifegate Impatto Zero). This crowdsourcing initiative was included because, when a non-profit and/or a government institution initiate crowdsourcing projects, they can trigger sense of civic duty [17], drive to contribute to the community and concerns about the democracy. These characteristics can be powerful motivators for individuals to participate [18] and they end up making the initiative community-driven.
- Almost all of the initiatives started after 2000.

As anticipated by the literature review, the picture of the mapped initiatives confirms an extremely diverse landscape of communities, associated by the presence of a clearly identified community goal and the essential role of web platforms.

3 Towards a Methodology for Online Initiative Assessment

The assessment of the impacts of community-based initiatives on society is the core goal of the TESS project, that developed three different methodologies to evaluate different results: a quali-quantitative assessment about initiatives' impacts; a qualitative assessment of their trajectories; and an assessment of the impact of the initiatives on the greenhouse gases (GHG) reduction.

The quali-quantitative assessment[5] aims to compare initiatives that are different in scope and structure and it focuses on four main areas of impact: social, political, economic, technological. The methodological framework identified relevant sub-dimensions for each area, and the appropriate indicators to evaluate the initiatives' performances[6]. The results are evaluated according to ten criteria for assessment, selected and ranked after extensive literature review and consultation with relevant stakeholders: Carbon reduction; Carbon efficiency; Social capital; Social inclusion; Political mobilization; External networking; Financial sustainability; Local economic impact; Innovativeness; Human capital.

The project also developed a qualitative methodology aiming to understand the up-scaling dynamics and the success factors of the initiatives, their evolution through time and how internal and external relationships, values and choices framed their activities and results[7].

Finally, TESS developed a methodology[8] to assess the impact of the initiatives under investigation on GHG emissions and reduction. The methodology is based on the GHG Protocol for Project accounting (WRI & WBCSD 2003), a framework developed by the World Resources Institute. The project adopted a comparative approach, calculating the emissions relative to a reference scenario which describes the hypothetical emissions in the absence of the activity though a detailed process including the identification of specific activities for each initiative and the identification of the appropriate baseline scenario and metrics for each country[9].

The analysis of online initiatives built on these three frameworks, adapting their research tools to develop a single questionnaire for two of the mapped initiatives, with the aim of gathering preliminary data about online initiatives and their performances and, when possible, about their differences and similarities with non-online initiatives. The descriptions emerged from the interviews do not allow a proper comparison or a full evaluation, but they trace first scenarios about online initiatives mechanisms and impacts and they make possible to identify questions for future research.

---

5 The original TESS data sheet for the multi-criteria analysis was elaborated by Filippo Celata, Venere Stefania Sanna and Cary Yungmee Hendrickson (Sapienza University of Rome).

6 D2.2: "Assessment data sheets for community-based initiatives", http://www.tess-transition.eu/wp-content/uploads/2015/07/TESS-Deliverable_2.2_FINAL.pdf.

7 D3.1: data collection database for success factors and constraints": http://www.tess-transition.eu/wp-content/uploads/2015/07/2015-07-01_TESS-Deliverable_3.1_FINAL.pdf.

8 http://www.tess-transition.eu/wp-content/uploads/2015/10/20151015_TESS_Del_2.1_final.pdf.

9 The TESS Track-It tool methodology has been developed considering country-specific dataset to calculate the CO2 reduction, according to the local context: http://www.sustainable-communities.eu/track-it/.

The resulting questionnaire is organized in two main sections: the first one aiming to assess the initiatives quali-quantitative performances, its goals, its perceptions and expectations, and the second one focusing on the financial assessment of the activities. The first part of the questionnaire is organized in five areas:

1. General information
2. Composition
3. Innovative effort
4. Benefits
5. Priority goals and activities

The questionnaire was sent to two initiatives and discussed during an online interview and further written exchanges. As illustrated in the following section, online initiatives did not provide enough data to fully exploit the methodology, but the questionnaire proved to be a useful tool to investigate them. The availability or the lack of data represents in some cases a significant result in itself, to draft conclusion and to address new research questions.

4 A Descriptive Analysis of Two Case Studies

Two initiatives were selected and interviewed to deliver this analysis. The selection of the two initiatives was based on the following criteria:

- Representativeness of different domains;
- Representativeness of initiatives originated in different countries;
- Representativeness of initiatives generated through different approaches.

The initial selection was potentially aiming to be larger, but the process was complicated by the difficulty in getting in touch with the initiatives or in getting their availability for the interviews: six initiatives were contacted with a negative outcome (meaning no answer or a negative answer), before reaching for the two addressed by this chapter: P2P Food Lab and European Cycling Challenge, both of which match the above mentioned criteria. These initiatives demonstrated good capacity to effectively summarize their story and to answer most of the questions.

4.1 European Cycling Challenge (ECC)

European Cycling Challenge (referred to as ECC, https://www.cyclingchallenge.eu) is a non-for-profit initiative that promotes a challenge among European cities: the city that, by engaging its citizens in using the bicycle, “rides” the longest total distance during the month of May wins. Participants subscribe to the ECC platform and download the challenge application on their mobile phones. During that month of May, the application registers the distance the users cover by bike and sum it up at city level. The city that manages to engage the higher number of citizens and the most committed ones (citizens riding more) achieves the best result. By gamifying the action of biking, the initiative offers a powerful tool to cities and citizens to promote the use of bicycles.

Cycling has emerged in recent literature not only as a sustainable form of transport, but also as both an asset and a result of a correct urban planning strategy. The use of the bicycle has been investigated in correlation with the main social, environmental and economic features of urban contexts, formalizing the idea of bikers as urban travelers and triggering several behavioral studies concerning their choices and attitudes [7]. Studying the strengths and the contradictions of this collective urban actor allowed the emergence of new debates about challenges and opportunities of their local context [1, 3] as well as their position within the physical activity and health debate [9]. ECC has proven to be a particularly successful project dedicated to cycling and bikers, because of its capacity to develop a sustainable model and to go beyond the end of the initial funding. It keeps growing over the years, as illustrated in the paragraphs below.

The ECC initiative pertains to the transport domain and it was initiated in Bologna, Italy. ECC was born as an experimentation of the European project CIVITAS MIMOSA, financed by the EC in 2010. MIMOSA involved five European cities exploring new approaches to sustainable transport. During the project, sixty volunteers in Bologna decided for the first time to trace the distance they covered by bike through their mobiles' GPS and find out who was riding more: the experiment was a great success given the strong appreciation declared by the participants. In 2011 the challenge was replicated involving seven European cities and enhancing the supporting technology, which started providing real-time ranking of the participant cities. The idea proved extremely successful and enticing and it kept running during the following years, empowered by the fast spreading of smartphones. The development of the ECC application made the participation easier for the bikers and the management easier for the organizers. According to the ECC Rules displayed on the initiative website, it is still possible to join the challenge without having a smartphone, by inserting data manually on one's own account. Apart from that and from the obligation to join one single city team, every kind of bike and of motivation to use it are valid.

The initiative is run by two people working full time during the months of April, May and June, and part time during the rest of the year, with a very limited engagement during Autumn and Winter, when they are likely to work on the ECC one day per week. The three main skills needed to run the initiative are international project management, ICT skills to supervise the platform and yearly improve it and to be a biker and possibly a participant to the ECC itself.

In May 2016 the challenge reached 46.000 bikers in 52 cities from 17 different countries, for a total of almost 4.000.000 km cycled in one month and with the winner city riding more than 800.000 km. Following comparisons revealed how the online nature of the initiative allows it to gather a number of participants higher than the average number of participants of the other CBIs analyzed by the TESS project and, at the same time, the actual action - and related impacts - are still at local level. On average, each participant rode around 90 km during the challenge. Countries that scored higher during the last edition were Polish, Swedish, Italian and French. Entering this result into the TESS Track-It[10] tool to generate an estimate of the potential CO2

[10] Given the ECC multi-country nature, the country selected to calculate the reduction was the one who performed the highest score during the 2016 challenge, which is Poland.

reduction of the initiative, it emerges that the ECC could avoid more than 500.000 kg CO2e in 2016, would the involved cities replicate their monthly performances during the entire year. This amount is equivalent to the emissions produced by 60 European citizens on average each year or to almost 2.300.000 km driven by car.

From the financial point of view, the initiative is no-for-profit and completely self-sustained: the only condition for cities to join the challenge is in fact a fee of 1000 €, that allow ECC to keep running and to develop each year the communication concept and the graphic. Then each city develops its own communication campaign around the participation to the challenge, involving the appropriate stakeholder and shaping the initiative to better suit the local approach to biking and sustainable transport. ECC usually recommends cities to join the challenge if they have enough time to adequately promote it or if they already have a strong biking community at the local level, to avoid wasting this opportunity: the simple participation without a consistent institutional support is not enough to reach the initiative final goal, which is biking promotion. During the years it has been proved that local community of bikers can actively require their cities participation to the ECC after hearing of it, as happened in the cities of Rome and Naples.

The main advantage that cities receive from their participation is the opportunity to enforce their own campaigns and communication activities through an international message. The second very relevant outcome is the data gathering and the traceability of bikers' itineraries: all bikers in fact register to the platform that will follow their movements for the entire month of May through the app on their mobile phones. According to the interviewee, these data, that cities get for free, would otherwise have a significant economic value. They allow cities to get a real-time picture of bikers' urban mobility and to frame their policies and infrastructural choices accordingly. They also allow tracing a profile of the city bikers, since at the moment of the registration they are asked about their year of birth, sex, zip code and employment. This represents a further incentive for bikers who, by participating to the challenge, can feel their direct involvement to the bike development of their own city. It is not possible to verify whether participants know each other before registering to the ECC, but according to the organizers it is likely that most of them do not have any previous relationship, thus confirming the crowdsourcing nature of the initiatives and the capability of activating a number of users that, even if present at local level, are not organized as a stable community. It would be interesting to understand if the participation in the ECC foster the reinforcement of local initiatives and/or the creation of new ones. According to the initiative's promoters' perception, one of the main strength of ECC relies in its trans-national dimension, which enlarges the pool of potential participants and reinforces the appeal of the challenge.

4.2 P2P Food Lab

P2P Food Lab (https://p2pfoodlab.net/index.en.html) is a non-for-profit initiative aiming at promoting sustainable lifestyles and at studying how technology can provide new opportunities for the local production, distribution and composting of food at community-based level. As all Peer-to-Peer initiatives, the final goal is to allow people to develop collaborative relationship to manage their own resources: the solution offered by the initiative is a platform to facilitate the production of food and the sharing of information. The first objective is achieved by sending a selection of seeds for participants to grow their own vegetables, and the second one by allowing them to share data about the process (weather, progress, photos) on the platform. In the long term, this allows the collection of a relevant dataset of information for people willing to join the sustainable agriculture community and it improves the opportunities for collaboration and innovative cross-fertilization among ideas, experiments and solutions.

The P2P Food Lab quickly reached the scientific debate thanks to the active engagement of its developers, positioning itself across the topics of the agro-ecology, open innovation and knowledge-based economy. Its development involved a highly interdisciplinary team made of engineers, scientists, citizens, artists, aiming to investigate open systems innovation for ecosystems leveraging and to build a wide-range collaborative platform starting from an e-laboratory [4]. The project built on the engagement of citizens and on the opportunities offered by Citizens Science for an innovative, bottom-up approach to sustainable food systems [5]. The project also built on previous studies on Alternative Farming Network and on the need for enterprises and society to move towards a more knowledge-based economy, with a focus on household gardening and the social, environmental and economic reasons beyond the phenomenon [11, 13].

P2P Food Lab started in 2012 in Paris as a project of the Sony Computer Science Lab, which is a research institute funded by the Sony corporation with the aim to investigate how technology can support sustainable and fair societal progress. The idea for this specific initiative arose from an exchange with a community in Bruxelles working on rooftop gardens-permaculture, which started using technology to measure and maximize their results. This brought to the idea of a peer-to-peer community having easy access to food production and sharing data about this activity, in fact a Citizen Cyberscience project: it relies on participants to set up their own experiments, gather data and collectively seek the most effective strategies to grow food. The community can use both quantitative data from sensors and free annotations from participants to build this knowledge base. The average number of participants is 100 each year and there are no specific criteria required to join the initiative. It is not possible to know whether participants knew each other before joining the P2P Food Lab, but it is likely that most of them did not, once more confirming that online platform allow to reach non-exclusive communities of peers.

Today the initiative is still hosted and financially supported by the Sony CSL but it is independent in terms of management. While the development of the idea saw the involvement of a pool of researchers, artists and activists, today P2P Food Lab requires a lean organization, with one person managing the platform and taking care of sending the seeds to the people that register to participate. The three main skills required to run the initiative are knowledge and awareness of societal needs and challenges, communication skills to promote the initiative and interact with participants, and high technical ICT skills as well as agronomic skills, both requiring around five years of education and five to ten years of experience. The annual cost to run the initiative - not considering labor costs - is around 1000€, where 300€ is the cost of all the seeds provided to the participants and the rest of the expenses cover sensor devices, servers and ICT equipment. Participants receive 5 grams of seeds, which allow them to growth around 10 kg of vegetables, with a likely saving of 20–30€ compared to purchasing the products through traditional distribution channels according to the project manager and a reduction of 275 kg CO_{2e} according to the TESS Track-It tool. This result is aligned other results emerged from the TESS project, confirming that regardless of how participants, beneficiaries or users come together as a community the emission reduction potential of food production tend to be low.

4.3 A Comparative Analysis

These final data makes clear the difference between the two interviewed initiatives, which are not comparable in terms of origin, activities and size. On the other hand, they share the overall goal to contribute to the development of a more sustainable world offering a collective space for communities of people brought together by a shared goal or lifestyle approach. While ECC activities are more oriented towards the promotion of values and awareness among the sensitive topic of urban mobility and planning, P2P Food Lab is, on a smaller scale, directly involved in the provision of goods and knowledge to shift part of the production system. For these initiatives the strong ethical motivation, previously identified as a shared feature among online initiatives and as a common marker with non-online initiatives, clearly emerged from both case histories. Regarding their contribution to CO_{2e} yearly reduction, the initiatives position themselves on the two extremes of the results identified by TESS main database for their field of activity[11]. About the food domain, P2PFoodLab is in fact aligned with small-scale initiatives generating low level of reduction, well below the "Provision of food" average of 2.151,8 kg CO_{2e} per year, while ECC results are far higher than the average reduction (25.134,38 kg CO_{2e}) for the "Transport of persons" domain. Results about more online initiatives would be necessary to compare online and non-online initiatives achievements.

The most relevant feature both initiatives have in common though is the management of the online platform, showing how very small teams can initiate and support wide and active communities. This also implies the possibility to reach a large group of people and possibly a relevant goal through a minimum and well-focused effort in terms of human and economic resources. Particularly relevant is the use of a gamification approach from ECC, which can be considered a peculiarity of online initiatives generally not applied by non-online ones, with a high potential for crowdsourcing participation among communities.

About the involved communities, ECC and P2PFood Lab present two separate models of community aggregation: in the first case, in fact, ECC draws its members from already established communities of cyclers, while P2PFoodLab crowdsources them from the public at large. For more in-depth investigation about those communities, the assessment methodology should be further refined about their composition and identity. The very high variance of the number of online initiatives members (Table 1) would require a segmentation of said initiatives according to the nature and

[11] D2.4 "Carbon reduction and community impact Scoreboard"(http://www.tess-transition.eu/wp-content/uploads/2016/06/2016_06_20_Deliverable_2.4_FINAL.pdf).

size and a customized analysis for the different categories. This point, together with other questions raised by this preliminary analysis, inform the following and conclusive section about the next steps for researching online initiatives.

Table 1. Selection of Online Initiatives

Initiative number	Name of initiative	Country of origin	Domain	Starting year	Members	Domain
1	Foodsharing.de	Germany	Food	2013	15.365	Food (5)
2	Ifoodshare	Italy	Food	2013	No data	
3	P2PFoodLab	Belgium (International)	Food	2012	70)	
4	Shareyourmeal	Netherlands (International)	Food	2012	45.000	
5	WWOOF	U.K. (International)	Food	1971	No data	
6	European Cycling Challenge	Italy (Europe)	Transport	2012	26.000	Transport (2)
7	Roadsharing	Italy (Europe)	Transport	2008	No data	
8	Earth hour	Singapore (Worldwide)	Energy	2007	7000 cities	Energy (4)
9	M'illumino di meno	Italy	Energy	2004	No data	
10	Lifegate Impatto Zero	Italy	Energy	2000	No Data	
11	Atmosfair	Germany (Europe)	Energy	2005	No data	
12	Freecycle	U.S.A. (International)	Waste	2003	9.134.953	Waste (4)
13	Intercanvis	Spain	Waste	2009	38 (not sure)	
14	Streetbank	U.K. (International)	Waste	2010	72.730	
15	Freegle	U.K.	Waste	2009	1.500.000	

5 Final Considerations

The literature review introducing this paper and the descriptive analysis of the two case studies allow to outline preliminary observations and to make some comparisons, identifying significant investigation areas.

Similarly to face-to-face community base initiatives, online initiatives offer a wide diversity of structures, purposes, dimensions and activities. The most distinctive traits in common with community-based initiatives in general are the aggregation of their members around a strong-shared goal and the need to compensate for something that is currently not available at societal level. Given their grassroots nature, online initiatives seem to have a strong tendency towards financial independency for their growth and persistence, and lack of strong endorsement from institutional actors, even if their emergence tend to be more often related to actors such as institutions (the European Commission for ECC) or foundations and research institutes with a private origin (the Sony Lab for P2P). This observation already emerged when describing the mapping done for this paper, requires further research to be properly explored and explained.

This research offers the opportunity to understand other key aspects shaping to online community based initiatives, which are worth of further investigation:

- How the use of ICT tools and platforms impact on these initiatives persistence compared to face-to-face initiatives, and if it influences their sustainability and their potential growth over time;
- If there are specific features of ICT tools that impact on the initiatives emergence and evolution;
- How initiatives generated online, which are currently connecting single citizens with common goals and needs, can reinforce local communities and/or create new local communities and how online initiatives users interact among themselves;
- If the management structure of the online initiatives appears similar in terms of members' participation to the decision-making process to face-to-face ones and if and how platform users are able or interested to influence the initiative management;
- The role of ICT, its potential to enhance already existing dynamics and activities or to bring out fully innovative solutions, revolutionizing the way in which communities jointly perceive themselves.

These dynamics allow understanding the relevance and the potential of investigating the role of ICT tools and platforms within community based initiatives, a role that is rapidly becoming more and more important. Online initiatives are expected to increase in number over the next years, and many face-to-face initiatives are progressively integrating in their management ICT tools, which are becoming essential to their own persistence. This transition phase is worth a full investigation, along the research questions outlined by this preliminary research.

References

1. Cupples, J., Ridley, E.: Towards a heterogeneous environmental responsibility: sustainability and cycling fundamentalism. Area R. Geogr. Soc. **45**(3), 283–291 (2013). (With the Institute of British Geographers)
2. Estellés-Arolas, E., González-Ladrón-de-Guevara, F.: Towards an integrated crowdsourcing definition. J. Inf. Sci. **XX**(X), 1–14 (2012)
3. Fischer, A., Peters, V., Neebe, M., Vávra, J., Kriel, A., Lapka, M., Megyesi, B.: Climate change? No, wise resource use is the issue: social representations of energy, climate change and the future. Environ. Policy Gov. **22**(3), 161–176 (2012)
4. Freuden-Pedersen, M.: Cyclists as part of the city's organism: structural stories on cycling in Copenhagen. City Soc. **27**(1), 30–50 (2015)
5. Funabashi, M., Hanappe, P., Isozaki, T., Maes, A., Sasaki, T., Steels, L., Yoshida, K.: Foundation of CS-DC e-laboratory: open systems exploration for ecosystems leveraging. In: Parrend, P., Bourgine, P., Collet, P. (eds.) First Complex Systems Digital Campus World E-Conference 2015. SPC, pp. 351–374. Springer, Cham (2017). doi:10.1007/978-3-319-45901-1_37
6. Hanappe, P., Dunlop, R., Maes, A., Steels, L., Duval, N.: Agroecology: a fertile field for human computation. Hum. Comput. **1**(1), 1–9 (2016). Pavlic & Michelucci

7. Leimeister, J.M., Sidiras, P., Krcmar, H.: Exploring success factors of virtual communities: the perspectives of members and operators. J. Organ. Comput. Electron. Commer. (JoCEC) **16**(3–4), 277–298 (2006). Lawrence Earlbaum
8. Moran, D., MacLeod, M., Wall, V., Eory, V., McVittie, A., Barnes, A., Rees, R., Pajot, G., Matthews, R., Smith, P., Moxey, A.: Developing carbon budgets for UK agriculture, land-use, land-use change and forestry out to 2022. Clim. Change **105**, 529–553 (2011)
9. Parkins, J.: Cycling and Sustainability. Emerald Group Publishing Limited, Bingley (2012)
10. Peters, M., Fudge, S., Jackson, T.: Low Carbon Communities. Edward Elgar Publishing Limited, Cheltenham (2010)
11. Saelens, B.E.: Environmental correlates of walking and cycling: findings from the transportation, urban design, and planning literatures. Ann. Behav. Medicine **25**(2), 80–91 (2003)
12. Seyfang, G., Smith, A.: Grassroots innovations for sustainable development: towards a new research and policy agenda. Environ. Polit. **16**(4), 584–603 (2007)
13. Shupp, J., Sharp, J.: Exploring the social bases of home gardening. Agric. Hum. Values **29** (1), 93–105 (2012)
14. Slee, B., Whitfield, R., Whitfield, S.: Discourses of power: the development of small-scale hydropower in North-East Scotland. Rural Soc. **21**, 54–64 (2011)
15. Wolfert, J., Verdouw, C.N., Verloop, C.M., Beulens, A.J.M.: Organizing information integration in agri-food—a method based on a service-oriented architecture and living lab approach. Comput. Electron. Agric. **70**(2), 389–405 (2010). Special issue on Information and Communication Technologies in Bio and Earth Sciences
16. WGBU: World in Transition: A Social Contract for Sustainability (2011)
17. Bonabeau, E.: Decisions 2.0: The power of collective intelligence. Sloan Manag. Rev. **50**(2), 45–52 (2009)
18. Howe, J., Junker, L.: Let us play. Assoc. Now, Am. Soc. Assoc. Executives **4**, 26–29 (2008)

Diversity in FabLabs: Culture, Role Models and the Gendering of Making

Christian Voigt[1(✉)], Elisabeth Unterfrauner[1], and Roland Stelzer[2]

[1] Technology and Knowledge, Zentrum für Soziale Innovation, Vienna, Austria
{voigt,unterfrauner}@zsi.at
[2] Happylab GmbH, Vienna, Austria
roland.stelzer@happylab.at

Abstract. Diversity and inclusion in the technology sector is increasingly debated, specially in the context of equal opportunities for all and a shortage of experts in many tech related industries. The need to be more inclusive can refer to different age groups, people with diverse culturally and linguistically backgrounds or gender. All in all, ethnic, gender and socio-economic diversity is not yet at the forefront of fabrication laboratories (FabLabs) agendas for change. This paper aims to contribute to the discussion of diversity and inclusion by primarily elaborating gender relations in FabLabs and, to a lesser extent, discussing age and socio-economic conditions of makers. Our analysis is based on 39 interviews and the analysis of 55,450 data points extracted from the log files of 3d-printers, CNC milling machines, laser cutters and cutting plotters. This combination of qualitative and quantitative data reveals that, indeed, some machines are used more frequently by men or women. However, the main difference is in absolute numbers, i.e. women are not joining FabLabs for a variety of reasons ranging from uninviting cultures to the lack of role models in technology driven areas in general.

Keywords: Digital social innovation · Making · FabLab · Gender · Log file analysis · Inclusion

1 Introduction

"With every new technology, social power relations and thus gender relations are negotiated" [1]. It is a frequent observation that technology has a history of being male dominated and that culture, language and habits are male and often incompatible with femininity [2]. We see socialisation still reproducing gender stereotypes in toys, role models, examples in school books or characters in movies and books. Maker spaces, as a blend of novel technologies but also new practices in developing and producing (almost) everything [3], are confronted with the same male dominance due to their technological nature. The so-called *gender gap* is often illustrated referencing the substantially lower numbers of female makers to be found in fabrication labs (FabLabs). However, it would be overly simplistic to reduce FabLabs to a set of shared machines open to the public. There exist different notions on what making entails, how making differs from hacking or what defines a maker space beyond the machinery

I. Kompatsiaris et al. (Eds.): INSCI 2017, LNCS 10673, pp. 52–68, 2017.
https://doi.org/10.1007/978-3-319-70284-1_5

found in a physical place [4]. According to [1], FabLabs are places that encourage collaborative learning, knowledge sharing and changing the local economy for the better. The question we want to raise is whether FabLabs can also be places where gender stereotypes are counteracted and to a degree reversed, either by moving technologies into areas such as design, human-computer interaction or generally presenting making as part of a bigger concept around self-expression and self-empowerment. Maker spaces have the potential to provide an experience with technologies that differs from other technology driven areas. We consider these reflections timely because there is a history of failures of initiatives and campaigns pushing for higher participation of women in technology related jobs and activities [5] and this might be a time to push for different or more consistent efforts.

The paper is structured as follows.

- A brief introduction to related works highlights the role of culture, cascading gender gaps and the lack of role models.
- We then discuss data from a recent round of interviews exploring the maker movement as a whole, with a particular focus on gender relations.
- Next, we analyse machine log data from one particular large FabLab and compare the insights from both types of research, the more qualitative interviewing and the more quantitative log file analysis.
- In the conclusion, we argue that the lack of diversity un FabLabs is a societal problem, already starting in schools and families. FabLabs, however, are in a unique position to address this issue and make it part of their agenda for making a difference in how today's society operates.

2 Related Work: Gendered Artefacts, Gender Gaps and Culture

That gender relations influence the design of machines is shown in various gender and technology studies, demonstrating that construction, functionality and descriptions of technological artefacts are one-sidedly influenced by male culture [6]. Still, an artefact's gender is not put in stone. A classic example is the microwave, which morphed from a bachelor targeting high-tech gadget into "a female-connoted kitchen appliance" [1].

However, it's not only a desire to get better products when policies aim to increase the share of women in STEM subjects (science, technology, engineering and mathematics). It is also about moving towards a fairer world where women have equal access to high tech jobs and career opportunities. Hence, laying eyes on the numbers only, is potentially neglecting the real challenge, which consists of getting rid of prejudices and misconceptions associated with gender and acknowledge areas where additional action is needed. Equal skills do not automatically lead to equal opportunities, as shown in studies on gaps in pay and promotion [7, 8]. Neither are 'gender gaps' restricted to the tech world, and are also present in early childhood education or nursing [9, 10], although here a lack of male professionals is noted that receives much less attention than the lack of women in tech.

The maker movement itself already has a strong political agenda that aims to correct a divide between a few producing and many consuming countries. Frugal innovations and the circular economy are concepts trying to empower those with less 'fabrication power' to provide solutions for markets with less 'purchasing power' [11], FabLabs are important creators of opportunities by enabling low cost entrepreneurship [12]. The gender divide, referring to tasks, pay and promotion unduly influenced by gender, could also be addressed by FabLabs, e.g. showing that discriminating genders based on skills is unwarranted. However, for that we need a culture in FabLabs that acknowledges the importance of gender relations without assuming that the 'right culture' will emerge automatically. Henry [13] set up a feminist hackerspace, which at the time was taken as a separatist effort of not wanting to join the larger DIY community. Unfortunately, the debate ignored that the reason for a gender imbalance at maker spaces may be the perceived disrespect or denigration of female makers: " *If we aren't at hackerspaces, it isn't because we don't make things, don't code, or aren't technical enough. It's because men act like the space is theirs. Women face harassment ranging from assault to much milder, but more constant, come-ons and innuendos. Our geek cred is constantly challenged or belittled*" [13].

3 Research Objectives

This paper has a twofold objective:

(1) Elaborating the role of gender relations in maker activities. Are there stereotypes or gendered perceptions of the typical female or male maker? Are there different expectations each gender brings to a maker space?
(2) Exploring the different insights from qualitative interviews and quantitative analyses of machine log data. For example, to what degree do insights overlap, complement or contradict each other? Are there activities or use patterns that are dominated by a specific gender?

The first objective was mostly addressed by interviewing makers. A multi-case study of 10 different make spaces [14], ranging from community organised types of maker spaces to maker spaces that were extensions of universities or companies, provided the context for our interviews. The second objective was addressed through the analysis of one of the FabLabs anonymised machine logs, including demographics of members since 2010, machine use statistics since 2012 as well as the type of membership chosen by each member during this time. For example, demographics were available for 3,306 makers in three cities (Vienna, Salzburg, Berlin). Although analytics of activities were later restricted to 2,723 Viennese makers, since they represented the largest share of makers with activity logs going back to 2012. Similarly, a set of six machine groups (3d-printers, CNC milling machines, laser cutter, cutting plotter, PCB etching and a transfer press) produced more than 133,000 data points, of which we analysed 55,450 after selecting the subset for the Viennese Makers.

4 Interviewing Makers from Diverse Backgrounds

This chapter shows the analysis of interview data from an ongoing research project[1] in order to elaborate gender relations in maker activities. The objective here is to see whether the aspects we found in the literature (Chap. 2), were also prevalent in the respondents' statements. We start with a brief description of the interviews and the method used for analysing the transcripts. The remaining three sections present interpretations of verbatim quotes, grouped under three themes: culture, gendering and role models.

4.1 Interviewing and Coding Process

One of the central aims in qualitative research is to (a) develop hypotheses and support theory building and (b) classify empirical data [15]. Classifying in our context means to systematically order the material following defined classification rules that are informed by theory and are empirically solid in the sense that they have discrimination powers, i.e. an observation can be mapped against an established category without too much ambiguities, even though the operational definition of categories is a known challenge in qualitative research [4]. The interviews were conducted in the context of an EU research project investigating governance, collaboration and value creation in maker spaces [16]. Hence, gender relations were a sub-section of a larger qualitative analysis of these interviews. While there are multiple techniques for interpretation available in qualitative research, classification and structuration were selected as the most appropriate forms for interpreting the cases.

The original material for the analysis is based on the transcripts of 39 interviews, resulting in 576 pages of coded text sections. The coding process was supported through MAXQDA (Fig. 1). MAXQDA[2] is a content analysis tool in qualitative research, supporting coding across multiple files. Codes can be established a priori (deductive) or 'on the fly' (inductive). In this study, we used a deductive coding approach, where the codes are derived from prior hypotheses. The resulting coding tree allows then for easy extraction of test snippets that have been associated with a particular code. The upper part of Fig. 1, for example, shows the set of documents being analysed and the lower part shows the beginning of the coding system, including sub-codes for main categories such as 'organizing' a makerspace and 'collaborative behaviours' in a makerspace. The final code system had 57 codes in total, but only one code referred explicitly to gender relations.

For the interviews, we chose 10 maker initiatives in 8 countries, following a purposeful sampling approach as described in [17]. The fundamental idea behind purposeful sampling is to make sure that categories of interest – in our case organisation, behaviour and value creation in maker spaces – are considered under as diverse conditions as possible. The exact sampling procedure would go beyond the scope of this paper, but we aimed to maximize variations in four dimensions:

[1] http://make-it.io.

[2] http://www.maxqda.com.

- degree of commercial orientation and, if applicable, business model;
- degree of democratization of decision making, e.g. open source principles;
- degree of inclusion of social values in the value propositions;
- scale of interactions at regional, national or global levels.

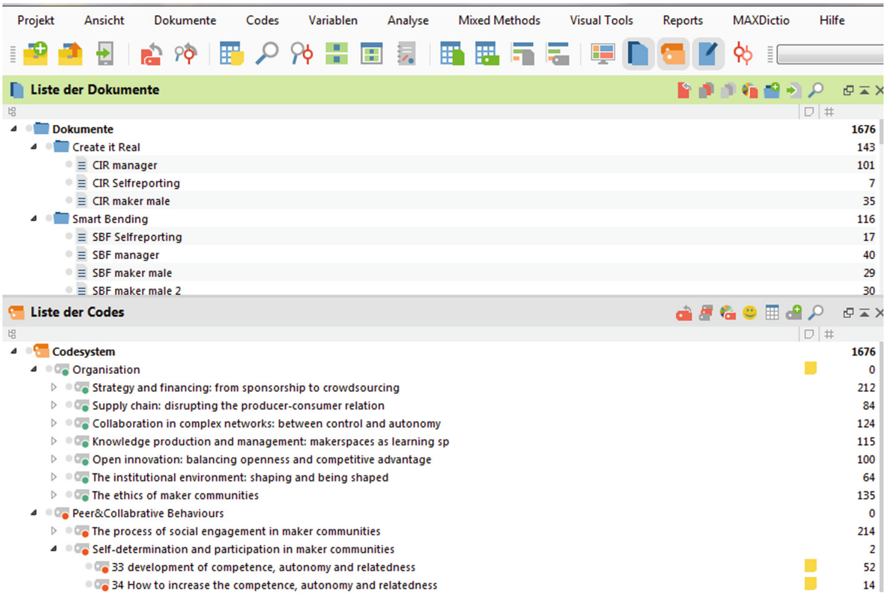

Fig. 1. Content analysis in MAXQDA (Source: Author's Screenshot)

The resulting maker initiatives included: Happylab Vienna, Austria (HAP); Danish Technology Institute Lab, Denmark (DTI); Fablab Barcelona, Spain (IAAC); Arduino, Italy (ARD); Hochschule Ruhr-West Lab, Germany (HRW); Dezentrale, Germany (Dezentrale); Mini Maker Faire Tartu, Estonia (AHHAA), Fablab Zagreb; Croatia (FLZ); Smart Bending Factory, the Netherlands (SBF); and Create It Real, Denmark (CIR).

The following three sections summarize interview responses related to gender relations. These sections will help to fine-tune future research directions under the headings of culture, gendering and role models.

4.2 Culture and Values

Culture is partly defined through what is valued. Or, put more succinctly, maker cultures are characterized by a 'hands-on' imperative, seeking to connect with people and one's environment in ways that allow to shape and change that environment [25]. People's attitude towards their environment is already formed very early on, when youths show different approaches to learning, taking things apart or integrating things [18].

In that context, one interviewee emphasized the need to actively promote a culture that appreciates diversity: *"And it depends of course also on who you admit and how you deal with mistakes, etc. You can quickly create an environment where making mistakes and also sometimes breaking something is not allowed. And by that you exclude people quickly"* (HRW Manager).

Then cultural influences continue in education where a "male-dominated and machine-fixated culture works to marginalize women" [19], directly influencing self-perception. As pointed out by Faulkner [5], the nerdy image of computer hackers and techies is seen as the antipode to the social engagement of women. Women are then forced to make a decision whether to remain within expected gender roles or not: *"But in these technical fields I think that boys are not better, but some girls are thinking 'OK, this is not a girl's job'"* (FLZ maker). A first step would be, to present gender relations as a serious issue: *"[The question of..] female presence in maker spaces hasn't been noticed much, now it's on our minds. For some time now we have started thinking in these terms—how many girls we have, etc."*[3] (IAAC manager).

Language in general seems to be an important issue, as shown in the following two quotes: "*When we speak of the Maker Faire and success stories to the media, I always try to keep in mind that we do not focus on the males only or we risk giving the impression it is an all-male gathering, which is completely false*" (AHAA Manager) and "*I also noted that a man talking to a woman about technology, he does it in a different way, like he has to explain it more thoroughly - they call it mansplain*" (ARD maker). Diversity should not be related to a minority of females in making or hacking. Dedicated events are needed, where female makers and hackers can be among themselves: *"Girlsday is always huge at our university, with about 150 female pupils. There, we always try that women or female university students facilitate the courses. That is the special thing about it. Due to the fact that they are among themselves, they are more relaxed"* (HRW maker). Similar spaces organised by women for women are Baltazar's Laboratory in Vienna and, in the past, also MzTek in London or Genderchangers in Rotterdam. These places are needed since most maker spaces give influence to those who can spend the most time in a space and "mark territory by gestures and male habitus, resulting in reduced comfort and diversity in a lab" [20]. Not the least it needs firm policies to avoid behavioural aberrations or investigate transgressions of codes of ethics or general norms of conduct: *"Well, we don't have posters of pin-up girls in our lockers or things like that and say: 'This is a man's closet and that a woman's.' or similar stories. We are aware of this. At the—[another maker space] they had a lot of these problems. The atmosphere there made women leave again"* (Dezentrale Manager).

4.3 'Gendered Task' and 'Gendered Objects'

Gender-based division of tasks can be commonly observed in today's society. Technologies in modern households are more often used by women than the technologies for home maintenance and gardening, which are predominantly used by men [5].

[3] Responses were corrected for grammar, however, we refrained from changing interviewees choice of words, hence the use of 'girls' and 'boys' with reference to makers.

In some instances, it's just a perception that implies relations between gender and technology*: "Women make bio stuff; men make 3D prints ... we also have women coming here now and then to 3D print and laser cut and so on, but the tendency is like that"* (Dezentrale manager). Or deterministic thinking about gender is expressed as a reference to past experiences: *"Most of the women were in communication, business and design. Now the one we have is doing design but at the same time, she can program"* (CIR Manager) or *"... [they] already employ female lorry drivers, but generally there are no women wanting to work in these factories"* (SBF Manager). Gendered activities are not only imposed from the outside - talking about activities - but is also used in describing, a person's own path of development: *"When I was younger I was doing more sewing and simple things for a female. I can't say I have moved to more technical things, but I have learned more topics, and possibly those sewing things are not so interesting anymore. Electronics is more interesting"* (AHHAA maker).

Gendered activities were also seen as instrumental in increasing female participation in making: *"What is nice about 3d-printing in schools, that you can print jewellery and girls become interested, and some other things for boys"* (FLZ manager) or *"Laser cutters are typical for girls, because 'My best friend is getting married and I can make a nice cake topper for her' - Pinterest plays a huge part in that... that's why maker spaces are booming"* (HAP, maker). In another case, the instrumentalisation of gendered activities did not work out: *"Actually we tried a stitching workshop but it didn't work"* (ARD, maker).

Just like activities, objects can be used in a gendered way as well [6]: *"in 8^{th} grade some girls printed a unicorn and the boys something like a cover for their mobile phone ... Many also look at what the others do and get inspired by that. If the neighbour says 'I make a mobile phone case', then they make the same thing"* (HRW maker). Ideally we can move towards perspectives where activities and objects do not have gender related connotations anymore and are expressed as straight forward problem solving experiences, as shown in this last quote: *"I don't remember where I got the idea. At some point I built this greenhouse, with the help of my grandfather, in the countryside. It's not very far away, but we don't go there every day, maybe weekly. It doesn't work well if it's always closed or open, and a couple of years after keeping this greenhouse open all the time, I decided I wanted to make it work properly, and I have the skills to do it"* (AHHAA maker).

4.4 Role Models

According to Kanter's theory on homo-social reproduction (cloning effects), people are more likely to find their ways into social environments if people with similar interests to their own are already part of the environment [20]. Learning from or working with female role models was also a strong theme in the interviews:

- *The problem that could be, is that if it's a man who is teaching, he would say it in one way, and a girl would say it in another way. They have role models, so if it is a woman saying it, so is it great to look up to her and who she is. It really depends on who is standing there.* (CIR maker)

- *I do think that by having the Fab Lab attended by a girl is sort of, it makes other women interested in coming to the Fab Lab.* (DTI Manager)
- *For the Girlsday we have looked for female students who implement it, as role models so to say.* (HRW maker)

As the few quotes above show, role models are needed in the most diverse situations (teaching maker skills, managing a maker space, organizing an event, etc.). Moreover, what is needed are 'success stories' able to capture the attention of potential makers, and 'success' doesn't need to be interpreted exclusively in terms of business: "*There are women, when they are there, who are usually very good and they get noticed quite a lot because they are special. We are setting up a partnership right now with a company where we are taking shares in the company, and the CEO is a woman. ... As you give good examples, as you bring good success stories on women using 3D printers, the end users will evolve"* (CIR Manager).

5 Analysis of Machine Log Files

Whereas the interview material originated from a variety of 10 cases, only the Happylab in Vienna had extensive log files, documenting the use of their machines. In some of the other cases such a system was not available because of a conscious decision to have a self-organizing community without too many regulations. The size of the community plays a role as well. Dunbar's number of 150 describes the size of as group with which we can maintain meaningful, reciprocal relationships, including trusting and relying on each other [21]. The Austrian FabLab whose data we analysed has currently about 2,200 members in three locations: Vienna, operating since November 2010, followed by Salzburg, opening doors in October 2014 and the Berlin facility, which started in September 2016. Hence the mere size of Happylab's membership was a strong motivator to adopt an automated approach to controlling who can access their machines, based on whether they had the required training or not.

5.1 Log File Curation

The log file analysis used data from 5 machine groups (number of data points are given in brackets):

- CNC Milling (7,004)
- Laser Cutter (37,275)
- Cutting Plotter (8,281)
- Printed Circuit Board Etching (526)
- 3D-Printing (2,364)

For the first four machine groups actual data on use duration could be analysed, 3d-printing was run through a different system automatically charging makers with different material costs. Hence, although we can establish frequency of 3D-printing and number of people doing 3d-printing, no data on the duration of print jobs were available. The comparatively high number of data points for laser cutters is due to the use of a dead man's switch, that required makers to confirm their presence at the machine every 3 min. In the first few months of operating the laser cutter, a new data point was created every 3 min, which was subsequently changed to only logging start and end times. This rather small detail highlights the importance of scanning even larger datasets for unusual patterns in order to detect and accommodate possible changes of measurement procedures. A further correction of the dataset included dropping point outliers, which can be caused by the faulty reading of a maker's ID card or because the maker didn't log out of the machine at all, until the next person used the machine. In this case we used 98% winsorization, effectively dropping all data below the 1st percentile and above the 99th percentile as outliers. Taking the distribution of spending per person as an example, cutting one percent of the data (4 data points) reduces the maximum of money spent from 7,000 € to 2,000 €, noting that the median of spending is 50 € per user in total (independent of his or her membership duration). Nonetheless we are aware that any type of outlier treatment introduces a statistical bias by either under- or overestimating the importance of an outlier. Alternative approaches are outlier classifications based on their median absolute deviation or a-priori thresholds. Hence, by using the percentile-based outlier test we could highlight some extreme values, which we then discussed with the FabLab managers. Figure 2 shows a density plot, where Kernel Density Estimates are calculated via the python package seaborn.distplot [22].

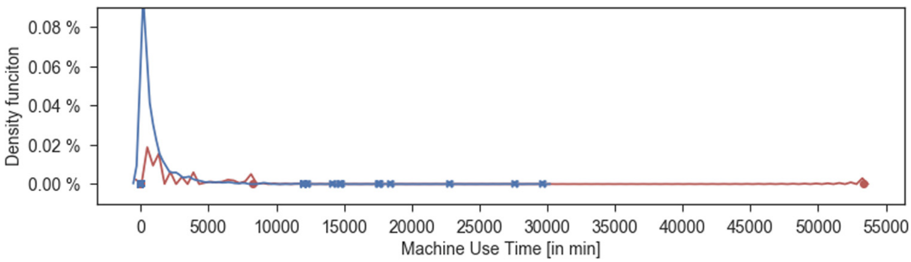

Fig. 2. One per cent outliers for laser cutting (blue crosses) and PCB etching (red dots) (Color figure online)

For example, power users could well amount up to 30,000 min in total for laser cutting (Median: 437 min, Interquartile Range: 144–1,124 min), whereas a value of 53,000 min for the etching printed circuit boards (PCB) seemed highly unlikely (Median: 210 min, Interquartile Range: 105–540 min).

5.2 Demographics of Membership

The following demographic data are restricted to the Viennese location of the FabLab, since maker activities will be discussed primarily in that specific location for the reasons already mentioned. The FabLab's membership is growing steadily until 2015 by more than 200 members each year. Having about 153 members in 2010, the FabLab had grown to more than 1627 members in 2016. Growth between 2015 and 2016 had flattened to 46 additional members. Of course, the objective for a physical location cannot be to continuously grow its membership since this could seriously affect the availability of machinery and induce longer waiting times for makers wanting to use a specific machine.

Looking at the percentages of female makers we can also see a steady growth until 2015, although at a lower rate than male makers. Since 2013 the percentage of female makers remains stable at around 26% (Fig. 3).

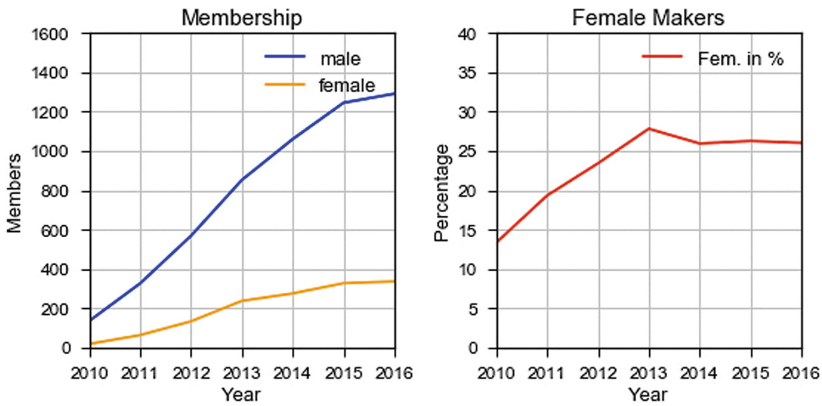

Fig. 3. Membership growth in absolute numbers and proportion of female makers [in %]

Looking at the age distribution of makers, we can see that male makers are slightly older (29 versus 32 years) and that female makers stay shorter with the FabLab (1.7 versus 2.4 years) (Fig. 4).

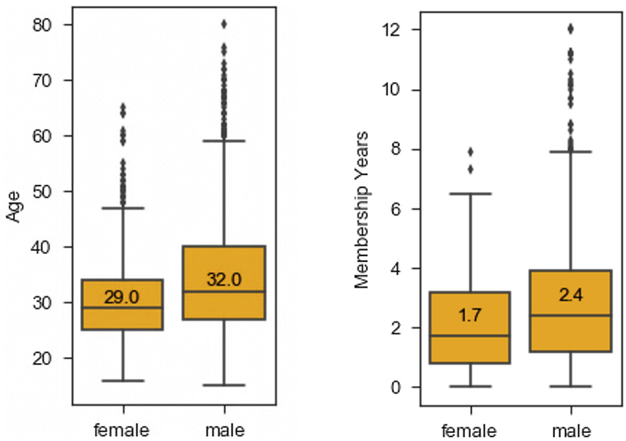

Fig. 4. Median age and membership duration (n = 2714)

Whereas the age difference might be negligible, the relatively small difference in median membership durations of 0.7 years could actually make a practical difference. Hypothetically, longer memberships in a FabLab could favour more diverse projects and more in-depth knowledge of what the machines can or cannot accomplish. However, further research is needed in order to make more conclusive statements about the influence of membership durations.

5.3 The Main Benefit of a Larger Membership is Flexibility

Beside the time someone is a member of a FabLab, another factor influencing his or her ability to actually get to know and use the machines is the type of membership a maker holds. Hence, we were interested in whether male and female makers show differences in their choices of membership types or in the extent to which they use the time available through their membership.

Concerning the Viennese FabLab, there are three types of memberships: (1) small - including three half days of four hours, (2) medium - like a small membership plus one extra day and (3) large - 24/7 access to the FabLab. Only looking at the makers who already had a track record in our machine log files, that is 542 male makers and 88 female makers (or 39% of the membership by the end of 2016), a snapshot shows that about 45% of male and 50% of female makers have chosen the small membership (Fig. 5), which is a moderate investment of nine Euros per month. The large membership has a 5.4 times higher price tag. Interestingly, the relatively more expensive 'large' membership is not far behind, with 42% and 38% of male and female makers respectively. No remarkable differences between genders can be seen at this point.

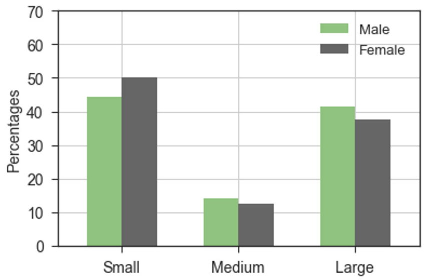

Fig. 5. Memberships in March 2017 (n = 630)

However, the split between the three membership types is constantly in flux, as makers can and do change memberships according to their needs. Looking at the time between 1.1.2012–31.3.2017 we analysed membership changes over 275 calendar weeks. Figure 6 shows two membership profiles over time to illustrate the issue. The profile in grey starts in week 30 and has frequent changes between a small and a medium membership and the green profile shows a primarily large membership with two breaks.

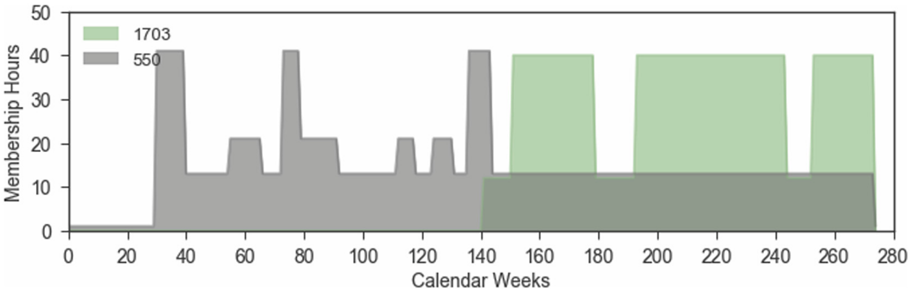

Fig. 6. Membership changes over time (Color figure online)

Nonetheless, the number of average weekly hours per person included in memberships remains relatively stable between 19 and 21 h. Figure 7 shows the development of the three membership types. Approximately every 52 weeks we see a dent in the graph, which reflects the terms and conditions of memberships allowing cancelation only at the end of the year (a condition now removed). Medium and large memberships could be cancelled or upgraded at any time.

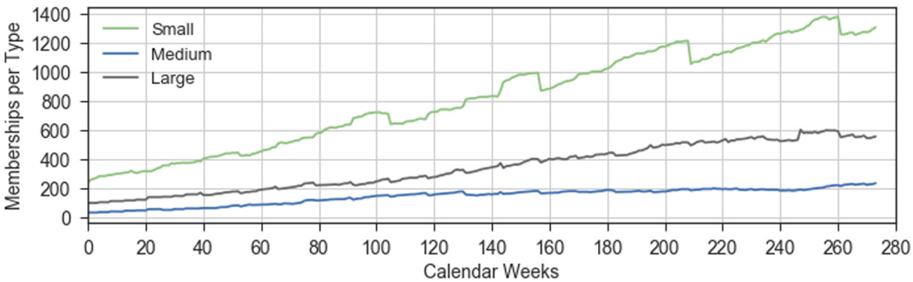

Fig. 7. Development of membership types (including all three locations)

We imagined that a large membership would also be associated with higher aggregated use times of machines. However, as shown in Fig. 8 (left side), there is no discernible pattern in the scatter plot, i.e. high use times can be found independently of the time included in a membership.

The right side of Fig. 8 confirms the independence of membership time and use rate, whereas ‘use rate’ is calculated as membership time in minutes divided by total minutes spent on FabLab machines. So far ‘use rate’ can only be taken as a calculatory proxy, because we do not have records of makers presence apart from the times they use a machine with access control. This implies that times on the electronics workbench, other manual tools or times for 3d-printing are not included. Nonetheless, keeping in mind these analytical limitations, we can see that the ‘use rate’ for the majority of makers is below 1%. Separating the genders, however, does not show a significant difference for 'use rates' between female and male makers: Mann-Whitney-U test (U = 38119, p = .35)[4].

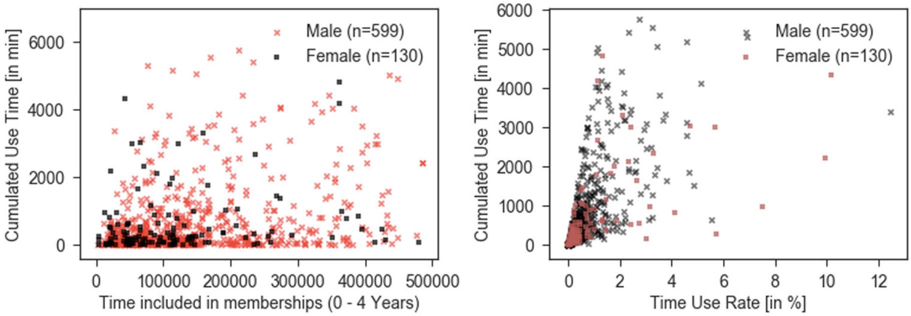

Fig. 8. An approximation to a time use rate

Given the relatively low use rates and the fact that more time included in memberships does not necessarily lead to more time used, we would suggest that makers opt for large memberships because of the flexibility that comes with 'large'. Hence, the main benefit of having access 24/7 is not the amount of time but the fact that the one hour a maker might need can be chosen at the most convenient time for the maker (Fig. 9).

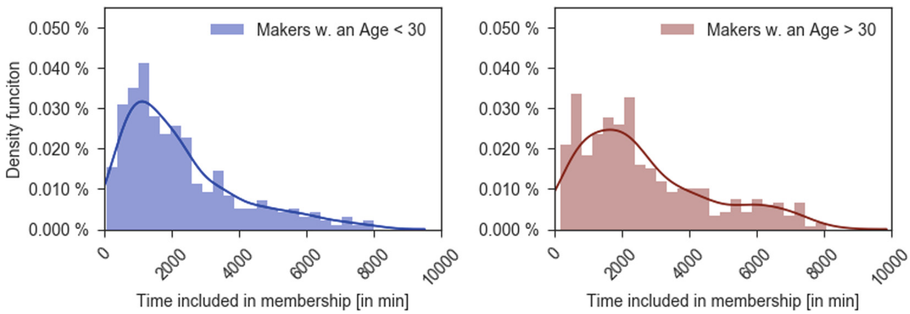

Fig. 9. Time included in memberships for makers below (n = 306) and above 30 years (n = 383)

[4] Since the machine log data were substantially skewed, i.e. not following a normal distribution, we used non-parametric statistics such as the Mann-Whitney-U test. Non-parametric methods do not require any a priori assumptions about parameters that describe the distribution of data (e.g. a mean or variance). The Mann-Whitney-U test is used when we want to know whether two samples (e.g. a variable split by gender) belong to the same distribution or not. Put differently, the test helps to establish whether there is a statistically significant difference between observations related to male or female makers.

Of course, this convenience has a price that might be more affordable to makers over 30 years. If we look at two samples of makers, we can see that makers over 30 have more time through their memberships (Median = 2,196 min) than maker under 30 years (Median = 1,744 min). Testing for group differences, we can say that the probability of a person over 30 years having more hours through their membership is statistically significant: Mann-Whitney-U test (U = 49122, p = .0004).

5.4 Gendered Preferences for Specific Machines

The next comparison of gender is related to the number of machines a maker uses as well as which machines are used. Of course, which machines are used depends on the maker's interests and projects, however, there is a relationship between knowing the machines and making informed choices about the most adequate way to fabricate a desired object.

Given a set of five machines (3d-printer, CNC milling, laser cutter, cutting plotter, PCB etching), 11.8% of female makers (n = 152) use more than two machines compared to 20.3% of male makers (n = 758). In terms of percentages, we can say that 3d-printer and CNC milling machines are more often used more by male makers, twice as many female makers as male makers use the cutting plotter and laser cutters as well as PCB etching are used to a similar degree by both genders (Fig. 10).

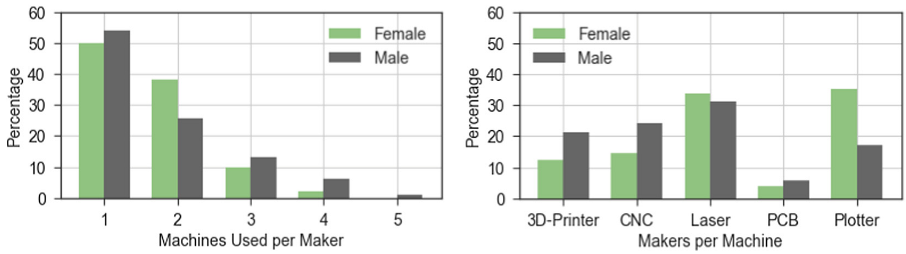

Fig. 10. Multiple machine use and percentages of makers per machine

In terms of median use time, we see moderate differences between genders: 15 min for cutting plotters, 34 min for CNC milling machines and the highest difference of 70 min can be found with laser cutters. The median use time for PCB etching was 78 min for male makers (n = 88). For female makers the overall number for using PCB etching was already quite low (n = 8) and restricting it to the first two years, only two female makers were counted (Fig. 11).

Further we also tested for the statistical significance of the observed differences, i.e. what are the chances that we see differences where there are none. Again, applying the non-parametric Mann-Whitney U test, all machines showed p-values substantially larger than .05, hence no significant differences between genders can be claimed.

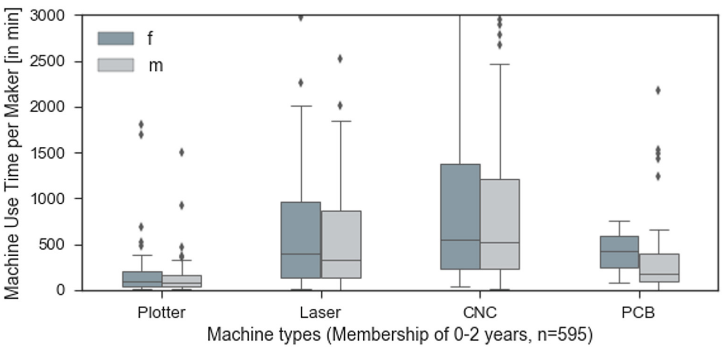

Fig. 11. Comparison of median use times based on machine type and gender

6 Conclusion: Gender Gaps and Lack of Diversity Are Not Specific to FabLab Activities

Summarizing the findings from the interviews, we have got statements indicating the importance of a respectful and supportive culture, the unwarranted genderisation of tasks and interests and the need for more female role models. The second part of the research was grounded in the logged activities of FabLab members and their demographic as well as membership data. The most overlap between the two data sets (interview transcripts and log data) was found in the area of analysing the actual behaviour of makers, at least their interactions with the machines.

Concerning a possible gender gap in FabLabs, the interviews revealed a strong link to early education and the wider presentation of women in tech. These are phenomena that originate largely outside the FabLab facilities, in families, kindergartens and schools. To a degree, we can see the same phenomenon reflected in the log data. The single most dominating difference are not varying use times but differences in absolute numbers between women and men joining a FabLab. There are four times more male makers than female makers in the first place, a gap that propagates into comparisons for specific machines, getting as high as 10 times more male makers than female makers trying printed circuit board etching. Other than that, both groups are relatively similar in age, female makers have slightly shorter memberships than male makers and are less prone to use 3d-printing or CNC milling machines. But then the two groups are comparable in their proportional preferences for the different membership types and the median use times spent with cutting plotters, CNC milling machines and laser cutters.

We are aware that a single FabLab is not representative for the whole maker movement, yet more than 3,000 makers and log data for about 4.5 years is a good starting point. Other FabLabs, such as the 'WeMake' FabLab in Milan, which is managed by a female maker, have gender ratios of 40% female makers versus 60% male makers (personal communication). However, the objective of studying gender relations or gender gaps within the maker movement cannot merely be to have equal shares of both genders. Hence, the debate should not get hijacked by surface numbers and more in-depth explorations are required, to better understand how to design spaces (physically and culturally) that are more attractive to a diverse range of people.

7 Outlook: Maker Spaces as Collective Awareness Platforms

Collective awareness platforms have been described by three trends: *social networking, pervasive computing (IoT, mobile platforms) and knowledge co-creation* [23]. To some degree these trends are already part of people's everyday life, be it because they connect through Facebook, use a smartphone including its diverse tracking sensors or consult Wikipedia. In principle, these trends are value free, meaning that 1.86 billion monthly active Facebook users use the platform to promote social good as well as hate postings against dissenting groups [24].

Claiming value neutrality for 'maker spaces' is not an option. Society's deeply engrained cultural stereotypes and value preferences when relating gender with technology, engineering or making do not stop at the doors of a maker space. As we have shown, prejudices do in parts persist and female makers have concerns about how they are perceived, accepted and whether they are granted the same opportunities as their male counterparts.

What we can see is a huge potential for technological action - without implying technology determinism - where communities of any gender, race or age can be empowered. That this won't happen automatically has been shown in various studies portraying the typical maker as white, male and wealthy - including extreme cases with non-white memberships of 0% and female members between 10–30% [25].

As outlined above, discussing gender requires more in depths reflection about gendered life-worlds (e.g. childcare obligations, less economic flexibility for part-time workers or simply interests, that do not match 'smart home applications'). FabLabs as collective awareness platforms do not necessarily need to offer childcare or e-stitching workshops [13], in order to attract female makers. As highlighted in section four of this paper, change starts with appropriate language, communicating equal expectations for girls and boys, while accepting different approaches to making (regardless of gender). At the end, collective awareness platforms stand for potential activities as in 'makers can…'. Aiming for more diverse maker spaces, we should avoid arguments including 'makers ought…' [cf. 26]. Moral exhortations will not get us the change we need, instead, what we need are physical and organisational designs of maker spaces that encourage the desired behaviours we want to see.

Acknowledgement. This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement 688241.

References

1. Carstensen, T., Walter-Herrmann, J., Büching, C.: Gendered fab labs. In: Fab Lab: Of Machines, Makers and Inventors, pp. 53–64 (2014)
2. Wajcman, J.: Feminism Confronts Technology. Polity, Cambridge (1991)
3. Gershenfeld, N.: How to make almost anything: the digital fabrication revolution. Foreign Aff. **91**, 43 (2012)

4. Voigt, C., Montero, C.S., Menichinelli, M.: An empirically informed taxonomy for the maker movement. In: Bagnoli, F., Satsiou, A., Stavrakakis, I., Nesi, P., Pacini, G., Welp, Y., Tiropanis, T., DiFranzo, D. (eds.) INSCI 2016. LNCS, vol. 9934, pp. 189–204. Springer, Cham (2016). doi:10.1007/978-3-319-45982-0_17
5. Faulkner, W.: The technology question in feminism: a view from feminist technology studies. In: Women's Studies International Forum, pp. 79–95. Elsevier (2001)
6. Berg, A.-J., Lie, M.: Feminism and constructivism: do artifacts have gender? Sci. Technol. Hum. Values **20**, 332–351 (1995)
7. Bishu, S.G., Alkadry, M.G.: A systematic review of the gender pay gap and factors that predict it. Adm. Soc. **49**, 65–104 (2016)
8. Blackaby, D., Booth, A.L., Frank, J.: Outside offers and the gender pay gap: empirical evidence from the UK academic labour market. Econ. J. **115**, F81–F107 (2005)
9. Mills, M., Martino, W., Lingard, B.: Attracting, recruiting and retaining male teachers: policy issues in the male teacher debate. Br. J. Sociol. Educ. **25**, 355–369 (2004)
10. Meadus, R.J.: Men in nursing: barriers to recruitment. In: Nursing Forum, pp. 5–12. Wiley Online Library (2000)
11. Tiwari, R., Fischer, L., Kalogerakis, K.: Frugal innovation in scholarly and social discourse: an assessment of trends and potential societal implications (2016)
12. Ames, M.G., Bardzell, J., Bardzell, S., Lindtner, S., Mellis, D.A., Rosner, D.K.: Making cultures: empowerment, participation, and democracy-or not? In: Proceedings of the Extended Abstracts of the 32nd Annual ACM Conference on Human Factors in Computing Systems, pp. 1087–1092. ACM (2014)
13. Henry, L.: The rise of feminist hackerspaces and how to make your own. Model View Cult. **2** (2014)
14. Millard, J., Deljanin, S.R., Sorivelle, M.N., Munk, K.B., Langley, D., Broek, T. van den Blankers, I., Voigt, C., Kieslinger, B., Unterfrauner, E., Pelka, B., Zirngieb, M.: D2.1 conceptual and methodological framework (2016). http://make-it.io/download/119/
15. Mayring, P.: Qualitative content analysis. Forum: Qual. Soc. Res. **1** (2000)
16. Unterfrauner, E., Voigt, C.: Makers' ambitions to do socially valuable things. Presented at the European Academy of Design Conference – EAD 2017, Rome, Italy (2017)
17. Palinkas, L.A., Horwitz, S.M., Green, C.A., Wisdom, J.P., Duan, N., Hoagwood, K.: Purposeful sampling for qualitative data collection and analysis in mixed method implementation research. Adm. Policy Mental Health Mental Health Serv. Res. **42**, 533–544 (2015)
18. Turkle, S., Papert, S.: Epistemological pluralism: styles and voices within the computer culture. Signs: J. Women Cult. Soc. **16**, 128–157 (1990)
19. Rasmussen, B., Håpnes, T.: Excluding women from the technologies of the future?: a case study of the culture of computer science. Futures **23**, 1107–1119 (1991)
20. Wuschitz, S.: Female Makers (2013). http://www.p-art-icipate.net/cms/female-makers/
21. Dunbar, R.: How many "friends" can you really have? IEEE Spectr. **48** (2011)
22. Waskom, M.: Seaborn Distribution Plot (Python Package) (2017). http://seaborn.pydata.org/tutorial/distributions.html
23. Sestini, F.: Collective awareness platforms: engines for sustainability and ethics. Technol. Soc. Mag. IEEE **31**, 54–62 (2012)
24. Zephoria: The Top 20 Valuable Facebook Statistics (2017). https://zephoria.com/top-15-valuable-facebook-statistics/. Accessed May 2017
25. Nascimento, S., Pólvora, A.: Maker cultures and the prospects for technological action. Sci. Eng. Ethics, 1–20 (2016)
26. Ozbekhan, H.: The triumph of technology: "can" Implies "ought". System Development Corporation (1967)

InSPIRES: Science Shops 2.0

Giovanna Pacini[1,3] and Franco Bagnoli[1,2,3](✉)

[1] Department of Physics and Astronomy and CSDC-Center for the Study of Complex Dynamics, University of Florence, Via G. Sansone 1, Sesto Fiorentino, Florence, Italy
{giovanna.pacini,franco.bagnoli}@unifi.it
[2] INFN, sez. Firenze, Florence, Italy
[3] Associazione Caffè-Scienza Firenze-Prato, Prato, Italy

Abstract. In this paper, we present some aspects of the European Project InSPIRES, which is focused on the concept of the science shop. A science shop is a methodology of opening universities and research centers to the civil society, accepting proposals for investigations and having them carried out by the students under the supervision and the scientific responsibility of a senior researcher. The goal of InSPIRES is that of cataloguing the methodologies used by the different science shops in the world, extracting the best practices; promoting the development of similar experiences in southern Europe and related countries, and increasing the continuous participation of citizens along the lines of the responsible research and innovation, also using Internet-based tools.

Keywords: Science shops · Science cafés · Responsible Research and Innovation · Science communication · Web 2.0 · Distant participation

1 Introduction

Science Shops [1] are one of the practical realization of the principle that the universities and research institutes should offer services directly addressed to citizens, clearly combining them with their fundamental mission of education, formation and research. This goal is a part of the *third mission* of these public institutions, funded at least in part by resources coming from the general taxation. In a certain sense, this engagement provides a *return* of the investments to the civil society and the socio-economic structures of the society. This return can be fulfilled in many ways: through the *technological transfer* (contract research and consulting with external clients, patents, creation of spin-off companies, participation in incubators and consortiums with transfer purposes), conservation services and dissemination, scientific and cultural training *outside the classroom* (management of archaeological sites and museum centers, popular conferences, lectures, articles and websites, participation in radio and television broadcasts), and so on.

The activities belonging to the third mission contribute, in Italy and other countries, to the evaluation of the universities and research centers and thus concur to the establishment of their financing quotes.

I. Kompatsiaris et al. (Eds.): INSCI 2017, LNCS 10673, pp. 69–78, 2017.
https://doi.org/10.1007/978-3-319-70284-1_6

Most of these activities, however, are either too specialized to reach the general public, that is unlikely to directly commission a research project at a university, or too unidirectional, i.e., not involving exponents of the *outside party* in a debate *among peers*.

Conversely, the engagement of the civil society is a strong request that comes from the European Union, that, through the Commission, developed the concept of *RRI*, which stands for Responsible Research and Innovation [2, 3].

As shown in Fig. 1, the concept of RRI includes actors such as researchers, the political bodies, the communities devoted to training and education, the industrial and business fabric, but also the civil society and its organizations.

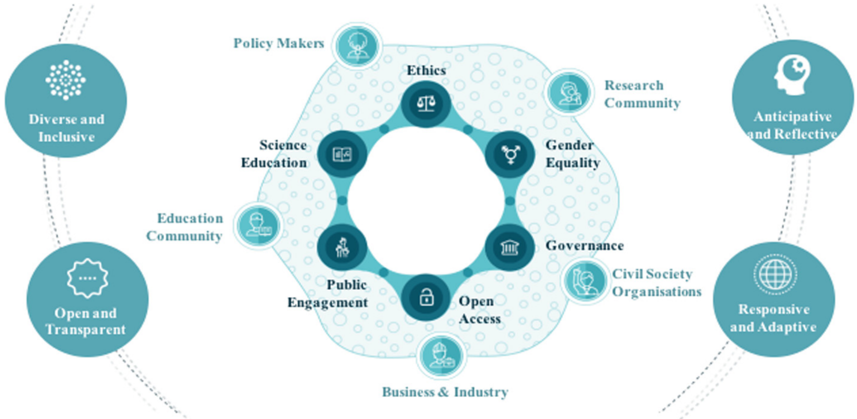

Fig. 1. A chart representing the pillars of Responsible Research and Innovation actions, image from Ref. [1].

The pillars of the Responsible Research and Innovation are:

- The involvement of the public from the beginning of any process (for example, a research) and during its course.
- Science education, in the sense of addressing a problem with the scientific attitude, without taking anything for granted and avoiding the principle of authority.
- Being aware of gender issues, and avoiding the discrimination, even subtle, based on sex.
- Considering the ethical issues and anticipating the problems with it.
- Ensuring free access to knowledge and to the results of scientific researches.
- Making the entire management process (governance) transparent and accessible.

This view was refined in all projects that are part of Horizon 2020, and especially in the Science with and for Society program [4].

Science shops [5–7] are a specific response to part of these requirements, through the provision, by universities and research centers, to undertake research proposals coming from the organization and the associations of the civil society.

2 Science Shops 1.0

We shall describe what might be called the version 1.0 of this practice, which was born in the Netherlands in the 70s [6], and then spread in Northern Europe and Canada, with examples in Austria, Romania and Portugal [5]. The concept of the Science Shop has been actively supported by the European Union, through various projects [8].

From the public's point of view, a science shop consists of an office, a series of meetings and a website (or other social network contacts) through which the civil society organizations and associations can submit questions about their problems. The questions can be of the most varied, from the impact study of a city change (closure of a road, opening of a railway line), management problems (social allotments threatened by the construction of new housing), health and environmental problems (pollution, stigma related to certain diseases), problems that have to do with the social perception of a topic (migrants, perception of risks) and so on.

The received question is analysed by the office in charge to decide whether it is ethically acceptable, and if the receiver has available human resources and suitable competencies to formulate a response. If so, the investigation is assigned to a researcher for its management and coordination, and a *call* is simultaneously open to students, who are the real *scientists*.

Students contribute to the research as part of the experimental work of a course, such as for example, courses in management or about planning surveys, or as their traineeship task or as their work for the final bachelor or master thesis. In some cases, universities provide a salary for students' traineeships, in other cases, the "reward" is constituted by the credits. In general, however, students consider it a very positive experience [9].

The theme of the research is refined and clarified in meetings with the proponent association, and the investigation results are presented and discussed in public meetings, in addition to being described in one or more publications. See the workflow in Fig. 2 (left).

This is only a coarse-grained description of the operation of a science shop, and there are many variations, for example in some cases the research is not managed directly by universities and research bodies, but it is rather carried out by specialized enterprises or spin-offs. There are also science shops specialized in very specific fields such as the design of bicycles or investigations related to rural problems.

This scheme works, but is fairly absent in southern Europe and related countries (for example the Mediterranean and Latin American countries), it is not particularly optimized and above all in principle it would allow much stronger interactions between the research institutions and the civil society, in the spirit of RRI.

Finally, the possibilities offered by the Internet are not particularly exploited. The InSPIRES project, which is described below, aims to explore some possible improvements and extensions of the science shop concept, and in particular has the ambition to improve the *one-way* flow of activities that, apart from the initial request and eventual meetings, sees all the initiative left to universities and research centres, with little participation by citizens.

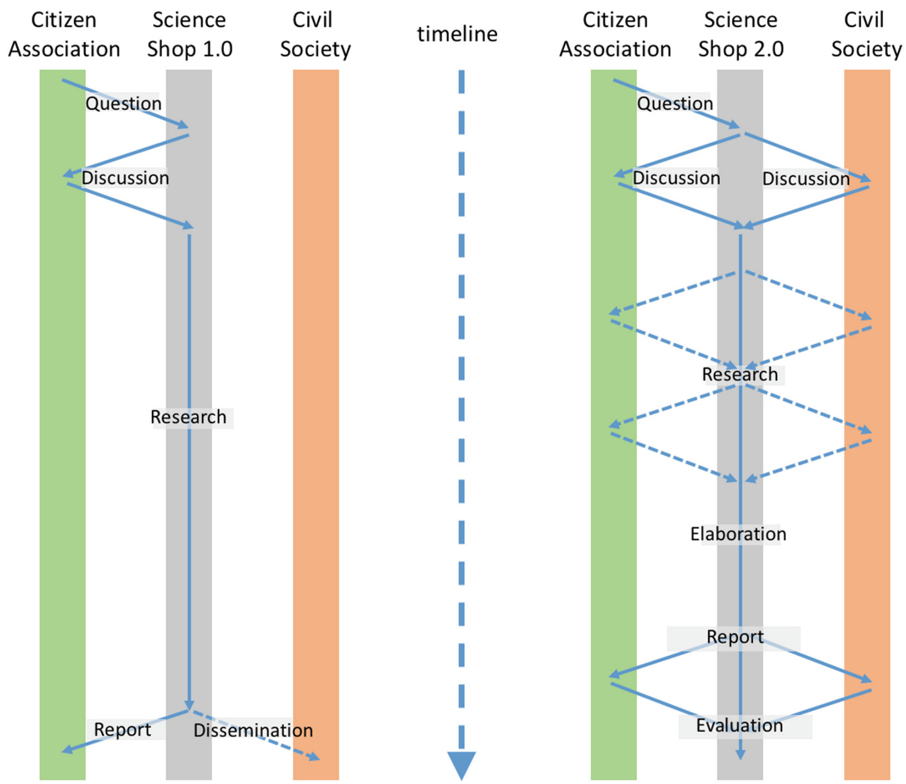

Fig. 2. A chart representing the (simplified) workflow of Science Shop 1.0 vs the one forecasted for Science Shop 2.0. Figure by the authors.

3 Science Cafés and Participation

A similar one-way problem in communication has led (1998) in the UK to the emergence of science cafés, inspired by the French philosophical cafés [10]. A science café is in some sense the opposite of a popular lecture.

In a conference there is a physical distance between the speaker/expert and the public, and also the choice of the place (a conference room) helps to ward off these two figures. In a certain sense, it is as if the conference takes place *in the expert's home*, and the audience was a guest.

As a result, the flow of information and interactions are almost always one-way, with a few questions from the audience confined to the end of the intervention. Participation is therefore very selective and are unlikely to involve a generic public. From the point of view of the interaction, the conference is still better than other forms of dissemination such as articles, television and radio broadcasts, but the latter reach a wider audience. Nowadays there are many instances of science cafés all around the world [11–13].

There have been many attempts to overcome this state of affairs, bringing discussions in places where people feel more "at home", such as science at the market [14], pint of science [15], science in the bar [16] and so on.

The science cafés are part of this initiative, but in principle, following their approach, the discussion itself is driven by the public, which must be *on par* with the experts, while respecting their expertise. An ideal science café (at least in the original British scheme) takes place in a pub, without slides nor microphone, and the discussion begins and is driven by questions from the audience.

With the increase of the public classic audio supports (microphone, amplifier and speakers) was added, as well as an audiovisual support (slides), so abhorred in the UK, that this mode is called the *french one*. The advent of the Internet has also allowed to use tools such as podcasts to distribute the event recording and also streaming, introduced mainly by our science café in Florence [17], as described in Ref. [18].

The *cognitive* aspects of the science cafés are the need to actively involve the public in the transmission of information, avoiding the *TV effect*, and simultaneously avoid taking premature position on *hot* topics, since there is the risk that a too large distance among the expressed points of view obstacles the exchange of information. Many of these aspects survive also when using *virtual* interactions [19–21].

The science cafés undoubtedly constitute a participatory mode of interaction, especially when they are supported by an active community. In the Florence case, the community consists of both academic and National (Italian) Research Council (CNR) researchers, but also *normal* citizens. The science café method will be one of the preferred participatory tool of the nascent science shop in Florence [22], motivated also by the fact that the *Caffé-Scienza* association of Florence and Prato [17] is part of the advisory board of the InSPIRES project.

4 Science Shop 2.0 and the InSPIRES Project

Our project has three main purposes: to highlight the *best practices* and excellences in running science shops, trying to catalogue the skills and methods developed in the previous projects and in the practice of present experiences; to promote the emergence of new science shops in southern Europe and related countries, and to try to develop a **2.0** mode, more interactive and in line with the RRI principles.

The first two goals imply an extensive review of the literature and interviews with a large fraction of administrators of existing science shops and of many subjects that would possibly give birth to new experiences, like rectors/directors of universities and research centres, bloggers and associations active in science popularization, citizen science activists [23], and so on.

The last point (science shop 2.0) will be faced trying to capitalize the knowledge about crowdsourcing, development of collective intelligence and participation developed and tested in other fields, such as that of the Collective Awareness Platforms for Sustainability and Social Innovation (CAPS) [24–27], and, concerning science cafés, the experience accumulated in SciCafe2.0 project [28].

Some methodologies, which will be tested shortly, have already been taken into consideration. First of all, in Florence we will test a *light* version of the science shop named *ask to UNIFI* [29], similar in some ways to the service offered by the network of public libraries of Lyon under the name *Guichet du Savoir* [30]. By means of this service, people can present simple questions how that can be directly *processed* by researchers through a simple bibliographic investigation, and receive a prompt response (perhaps accompanied by a small video. In this way we can ensure a quick return to the population, and stimulate the participation of our researchers.

In addition, we plan to use to systematically exploit Internet capacities, using the experience accumulated with the science café practice. Upon the acceptance of an applications will use tool such as mailing lists and forums to encourage citizens to participate in the analysis of the problem and propose generalizations, identifying related issues that could benefit from research, singling out issues related to privacy, ethics and issues gender and anticipating possible objections. Obviously the science café and other participatory discussion tools (world cafés, Delphi methods, ritual dissidents, etc.) will be used to extract knowledge from face to face and remote encounters. See the workflow in Fig. 2 (right).

5 The Internet Dimension

The “telematics” and social possibilities offered by the Internet are not systematically used by present science shop, except of course using the web for the promotion of initiatives and for stimulating the participation. For example the Lyon science shop has a pleasant and attractive site [31].

In general, however, citizens have little way of being directly involved in research or to participate in the definition of the problem, at least not through the Internet, which of course is justified by the fact that in general one tries to involve local communities in *real* interactions. Nowadays however this could also represent a point of weakness, given that many communities are dispersed or even completely virtual.

In particular, in the InSPIRES project we will also try to use methods that are derived from the *Citizen Science* [32] experiences to allow citizens to directly participate in the researches.

6 Citizen Science and the Direct Involvement of Citizens in Science Shops 2.0

If we look up the definition of citizen science on the Oxford Dictionary we find the following sentence: *scientific work undertaken by members of the general public, often in collaboration with or under the direction of professional scientists and scientific institutions*. From an abstract point of view, it is a very attractive and democratic idea, perfectly in line with the RRI objectives. What could be better than to involve citizens as researchers in an investigation, as indeed is done with students in a science shop?

Unfortunately, then passing to the practice, the experiments which have been tested are not very exciting. For instance, in Florence we are carrying out an initiative that sees the Institute of Biometeorology (IBIMET) of CNR [33], the Department of Information Engineering of the University of Florence [34], a SME (Magenta LAB [35]) in collaboration with the CSDC/Department of Physics and Astronomy [36] work together to develop and use low-cost sensors for pollutants and traffic detection. These sensors are distributed to citizens, involving them also in a survey about their perception of pollution, and in parallel modelling the evolution of the traffic and the diffusion of pollutants using data from such sensors. But the involvement of citizens remains limited, especially for what concerns their creative contribution.

In many cases, problems submitted to science shop concern the opinion and perception of citizens. An example is the one discussed above about the perception of chemical pollution, to which one could add the *perceived* risk due to electromagnetic radiation. The public service responsible for monitoring these pollutants (here in Tuscany the ARPAT service [37]) are sometimes overloaded by requests arising from the **perception** of a risk rather than the evidence of its presence.

There are many other examples, for instance the perception of the risks related to migrations such as crime, safety and transmittable diseases, which contrasts with the risk of being stigmatized (and the avoidance of diagnostic tests) if one is recognized as a carrier of diseases such as HIV or Chagas. Other examples of researches in which we have directly participated, and that could be suitable subject of study in a science shop, regard the risk and spread of self-injury in adolescents [38], the prevalence of Intimate Partner Violence (IPV) in couples [39] and so on.

In these cases, the investigation is carried out using questionnaires and trying to involve as many actors as possible. However, each questionnaire requires the identification of the participant from the point of view of his/her social, psychological and economic profile, as well as collecting standard data concerning for example the gender. A large part of any survey is dedicated to such profiling, and indeed our experience with psychologists tells that they would love to know many aspects of participants, countered by the fact that a too long profiling phase discourages and inhibits participation.

A possible solution would be that of actively involving citizens as long-term *sensors*, developing a web or mobile application that allows to collect their profile only once, storing it in an anonymous form on a server, so that the user only have to identify him/herself, eventually updating his/her own profile, and then answer to the *question of the day*. Of course, the citizens would have full control over their data, guaranteed by the public role of the University, and may indeed use this interface for commenting and deepening the qualitative and emotional/sentimental aspects related to the topic under investigation and also for proposing other topics.

In fact, similar needs are certainly not limited to problems that arise during the science shop activities, but are common to many other European projects, at least that involving interactions with citizens, for example, many projects in the CAPS and Science with and for Society calls.

Obviously we do not think we can develop such a service within a small project like InSPIRES. It would be desirable if such a problem were developed at the European level, which would ensure compliance with the ethical and privacy rules of the European Union. Indeed, in many cases similar profiles are collected anyway level by individual projects, a centralized management would guarantee a higher level of privacy and security. It would be a *responsible* (in the sense of RRI) version of the collection of citizens' data that is done in a very opaque way by many private services such as Google, Facebook and the like.

7 Conclusions

We have presented some of the guidelines that underpin the European InSPIRES project, involving the development of participative methodologies applied to the practice of Science Shops. These methods are inspired and mediated by other participatory experiences such as science cafés and citizen science, with which our project hopes to establish synergistic connections. We are also exploiting these connections for developing *telematics* and *social-networks-oriented* extensions for our science shops.

Acknowledgements. This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 741677 (InSPIRES).

References

1. European Commission: EUR 20877—Science Shops - Knowledge for the community (Luxembourg: Office for Official Publications of the European Communities 2003). ISBN 92-894-6246-9
2. RRI-Tools European Project web site (2017). https://www.rri-tools.eu/about-rri
3. Turk, Ž., Sessa, C., Morales, S., Dupont, A.: Responsible research and innovation in ICT – a framework. In: Bagnoli, F., Satsiou, A., Stavrakakis, I., Nesi, P., Pacini, G., Welp, Y., Tiropanis, T., DiFranzo, D. (eds.) INSCI 2016. LNCS, vol. 9934, pp. 239–243. Springer, Cham (2016). doi:10.1007/978-3-319-45982-0_21
4. Science with and for Society web site (2017). https://ec.europa.eu/programmes/horizon2020/en/h2020-section/science-and-society
5. Science Shop Wikipedia entry (2017). https://en.wikipedia.org/wiki/Science_shop
6. Leydesdorff, L., Van Den Besselaar, P.: What we have learned from the Amsterdam science shop. In: Blume, S., Bunders, J., Leydesdorff, L., Whitley, R. (eds.) The Social Direction of the Public Sciences. SOSC, vol. 11, pp. 135–160. Springer, Dordrecht (1987). doi:10.1007/978-94-009-3755-0_6
7. Living Knowledge International Network of Science Shops (2017). http://www.livingknowledge.org/
8. Science Shops European Projects. http://www.livingknowledge.org/projects/overview/eu-funded-projects/
9. Fokkink, A., Mulder, H.A.J.: Curriculum development through science shops. Environ. Manag. Eng. J. **3**, 549–560 (2004)
10. Science café, Wikipedia entry. https://en.wikipedia.org/wiki/Caf%C3%A9_Scientifique
11. British network of science cafés. http://www.cafescientifique.org/
12. US network of science cafés. http://www.sciencecafes.org/
13. Italian network of science cafés. http://www.caffescientifici.it/
14. Science at the market, University of Illinois. http://physics.illinois.edu/outreach/science-at-the-market/
15. Pint of science. https://pintofscience.co.uk/

16. Ciencia en el bar. http://www.cienciaenelbar.com/blog
17. Florence science café. http://www.caffescienza.it/
18. Bagnoli, F., Pacini, G.: Science Cafés in the internet era. In: Kompatsiaris, I., et al. (eds.) INSCI 2017. LNCS, vol. 10673, pp. 79–89. Springer, Cham (2017)
19. Duradoni, M., Bagnoli, F., Guazzini, A.: Non-trivial reputation effects on social decision making in virtual environment. In: Bagnoli, F., Satsiou, A., Stavrakakis, I., Nesi, P., Pacini, G., Welp, Y., Tiropanis, T., DiFranzo, D. (eds.) INSCI 2016. LNCS, vol. 9934, pp. 115–122. Springer, Cham (2016). doi:10.1007/978-3-319-45982-0_11
20. Guazzini, A., Cecchini, C., Guidi, E.: Small group processes on computer supported collaborative learning. In: Bagnoli, F., Satsiou, A., Stavrakakis, I., Nesi, P., Pacini, G., Welp, Y., Tiropanis, T., DiFranzo, D. (eds.) INSCI 2016. LNCS, vol. 9934, pp. 123–132. Springer, Cham (2016). doi:10.1007/978-3-319-45982-0_12
21. Perfumi, S.C., Cardelli, C., Bagnoli, F., Guazzini, A.: Conformity in virtual environments: a hybrid neurophysiological and psychosocial approach. In: Bagnoli, F., Satsiou, A., Stavrakakis, I., Nesi, P., Pacini, G., Welp, Y., Tiropanis, T., DiFranzo, D. (eds.) INSCI 2016. LNCS, vol. 9934, pp. 148–157. Springer, Cham (2016). doi:10.1007/978-3-319-45982-0_14
22. Florence Science Shop pilot project. http://science-shop.complexworld.net
23. Irwin, A.: Citizen Science: A Study of People, Expertise and Sustainable Development. Routledge, London and New York (1995)
24. Sestini, F.: Collective awareness platforms: engines for sustainability and ethics. IEEE Technol. Soc. Mag. **12**, 53 (2012). doi:10.1109/MTS.2012.2225457
25. Arniani, M., Badii, A., De Liddo, A., Georgi, A., Passani, A., Piccolo, L.S.G., Teli, M.: Collective Awareness Platforms for Sustainability and Social Innovation: An Introduction (2014). http://booksprints-for-ict-research.eu/wp-content/uploads/2014/07/BS5-CAPS-FIN-003.pdf
26. Pacini, G., Bagnoli, F.: Results of a collective awareness platforms investigation. In: Bagnoli, F., Satsiou, A., Stavrakakis, I., Nesi, P., Pacini, G., Welp, Y., Tiropanis, T., DiFranzo, D. (eds.) INSCI 2016. LNCS, vol. 9934, pp. 19–26. Springer, Cham (2016). doi:10.1007/978-3-319-45982-0_2
27. Bagnoli, F., Guazzini, A., Pacini, G., Stavrakakis, I., Kokolaki, E., Theodorakopoulos, G.: Cognitive structure of collective awareness platform. In: 2014 IEEE Eighth International Conference on Self-adaptive and Self-organizing Systems Workshops (SASOW), pp. 96–101. IEEE, London (2014). doi:10.1109/SASOW.2014.38
28. Badii, A., et al.: Collective awareness platforms and digital social innovation mediating consensus seeking in problem situations. In: Bagnoli, F., Satsiou, A., Stavrakakis, I., Nesi, P., Pacini, G., Welp, Y., Tiropanis, T., DiFranzo, D. (eds.) INSCI 2016. LNCS, vol. 9934, pp. 55–65. Springer, Cham (2016). doi:10.1007/978-3-319-45982-0_5
29. University of Florence (UNIFI) web site. http://www.unifi.it/
30. Guichet du Savoir. http://www.guichetdusavoir.org/
31. Boutique des Sciences de Lyon. http://boutiquedessciences.universite-lyon.fr/
32. Citizen science Wikipedia entry. https://it.wikipedia.org/wiki/Citizen_science
33. Biometeorology Institute (IBIMET) of the National (Italian) Research Council (CNR). http://www.ibimet.cnr.it/
34. Department of Information Engineering, University of Florence. http://www.dinfo.unifi.it/
35. Magenta Software Lab. http://www.magentalab.it/?lang=en
36. Centre for the Study of Complex Dynamics, University of Florence. http://www.csdc.unifi.it/
37. Agenzia Regionale per la Protezione Ambientale della Toscana (Tuscany Regional agency for Environmental Protection - ARPAT). http://www.arpat.toscana.it/

38. Remaschi, L., Cecchini, C., Meringolo, P.: Community-based strategy to prevent deliberate self-harm in adolescence: an inquiry to find risk factors at school. Int. J. High Risk Behav. Addict. **4**(1), e19663 (2015)
39. Guidi, E., Meringolo, P., Guazzini, A., Bagnoli, F.: Stochastic agent-based models of intimate partner violence. arXiv:1611.06544

Science Cafés in the Internet Era

Giovanna Pacini[1,3] and Franco Bagnoli[1,2,3](✉)

[1] Department of Physics and Astronomy and CSDC-Center for the Study of Complex Dynamics, University of Florence, Via G. Sansone 1, Sesto Fiorentino, Florence, Italy
{giovanna.pacini, franco.bagnoli}@unifi.it
[2] INFN, sez. Firenze, Florence, Italy
[3] Associazione Caffè-Scienza Firenze, Florence, Italy

Abstract. In this paper, we analyse some experiments of extending the practice of the science cafés, a well know participatory methodology of science communication, with the help of the possibilities offered by Internet. We describe the use of a web platform to disseminate the events and the use of YouTube for streaming and podcasting, with a comparison between costs and benefits.

We have experimented and analysed several scenarios, from the streaming of traditional science cafés to the dissemination of the results of European projects.

Keywords: Science café · Science communication · Web 2.0 · Distant participation

1 Introduction

The goal of this paper is that of illustrating some experiments about extending the participative methodologies of science cafés using some tools typical of the Web 2.0. In particular, we exploited a custom chat interface developed within the EU project SciCafe2.0.

The scheme of this paper is the following: In Sect. 2 we quickly review the motivations that led some scientists and science communicators (and us) to adopt the science café approach. In Sect. 3 we illustrate the composition of the science café public in four Italian cities and in Sect. 4 we describe our experiments about extending science cafés to the Internet. Conclusions are drawn in the last section.

2 The Evolution of Science Communication and the Birth of Science Cafés

Science communication can be schematically divided into institutional communication and popular dissemination, although there are many variations in between. The institutional scientific communication is primarily based on the presentation of the results of the researches in conferences that generally include a discussion of the results, followed by the publication of the scientific results in specialized journals, after a review phase. In most cases the presentation makes use of specific and precise terms and formalisms,

I. Kompatsiaris et al. (Eds.): INSCI 2017, LNCS 10673, pp. 79–89, 2017.
https://doi.org/10.1007/978-3-319-70284-1_7

that are easily understood only by specialists. What is more important, in any phase the dialogue is established among *peers*. A famous example of this approach is the Physical Review Letters rejection of a paper by Einstein and Rosen, after a negative review [1].

Conversely, popular science communication generally implies the disclosure of scientific results, without going into too much details, using everyday language and avoiding as much as possible the use of a specialized formalism. The special case of Wikipedia can serve as an illustration of the two approaches. Wikipedia is oriented toward a broad audience, but it is not written by specialists in popularization. It is not rare to find there articles written by specialists, which are not easily understandable even by scientists coming from a different field.

The two main differences that we want to stress between institutional communication and popularization is the *peer* vs. *top-down* approach and the *interactive* vs. *passive* role of the audience.

The two modalities of communication, institutional and popular, also differ in the media used. It is rare to find scientific communication conveyed by television, radio, websites, social networks and so on, while there are many broadcasts, websites or videos on YouTube with outreach purposes. The only exception to the rule is represented by the web sites of individual researchers or research groups, which normally report their scientific activity and sometimes offer the possibility to download the published work, seldom accompanied by a generalist introduction or explanatory texts.

With the increasing importance of science and technology in everyday life and the need of reaching an *informed consensus* concerning specific political, social or moral options (e.g., nuclear power, privacy, artificial insemination, genetically modifies organisms, etc.), both scientists and the public feel the need to participate in the scientific debate, or to be informed and to be able to express one's point of view.

This need was first realized in the '80s, and scientists were asked to address the broader public by communicating their results through various initiatives:

> *Scientists must learn to communicate with the public, be willing to do so, and indeed consider it their duty to do so. All scientists need, therefore, to learn about the media and their constraints and learn how to explain science simply, without jargon and without being condescending* [2].

This is what was written in the "Royal Society Report on the Public Understanding of Science" (PUS) of 1985, when a campaign for improving the public scientific literacy was launched in Britain.

Scientists were asked to take an active part in the public understanding of science and to consider it on of their duty [2–4]. This model, defined by the experts *deficit model*, has been used for years to interpret the relationship between science and society and involves a one-way communication from experts, scientists, towards non-experts [4–6]. In those years, in fact, existed the belief that:

> *The origin of the public controversy concerning the science or technology is the fact that citizens are missing an understanding of the knowledge, theories and scientific methodologies* [7].

In other words, scientists simply had to *translate* their scientific language into simpler words; this would settle the previous dispute. In this way, a relationship of trust between the public and scientific institutions that would have led to an extensive understanding of science was established.

This type of model is based on certain rules of basic communication:

- The model is one-way, from the scientist to the public;
- The scientist does not have a return of its communication;
- The public is not differentiated by education and the welfare state, and is regarded as *illiterate* from a scientific point of view [8].

After less than two decades, in 2000, the PUS was declared failed by the House of Lords; the analysis of the results showed that the level of literacy in England had not grown. The model that considered the interlocutors as a "unique" subject to which pass information, simply translated with a simplified language, failed.

> *It appears to be widely accepted that 'Public Understanding of Science' and 'Public Awareness' are too restrictive; 'Public Engagement' is now a standard reference, representing a commitment to dialogue. The larger framework of science and society now requires account to be taken of many complex dimensions of communication. The shift from 'public understanding' to 'public engagement', and the attention to the nuanced difference of 'science and society' and 'science in society' do reflect heightened awareness of the challenges of public communication of science* [9].

The *top-down model* was widely studied and several revisions have been suggested [10–16].

Officially PUS died in 2002 in favour of PEST, the *Public Engagement with Science and Technology*:

> *The Public Understanding of Science was no longer enough to engage today's opinions more sceptical and less deferential public. The phrase has a condescending, even demanding, tone which, so far from engaging in the public debate, tends to turn people off. In its place, we called for a new mood of dialogue: instead of the one-way, top-down process of seeking to increase to peoples' understanding of science, there has to be a two-way dialogue, where those seeking to communicate the wonders of their science, also listen to the concern of the public. Dialogue requires ears as well as voices* [17].

The science cafés practice anticipated this switch by decades. They were launched officially in 1996 with an article by Dallas [18].

A science café is a discussion about some topic of science and technology as scientists do. This means discussing with experts all on the same ground, where the public, and not the experts, are at home. This implies for instance not to run the discussion in a conference hall, but rather in a pub, with experts at the same level as the public.

The purpose of the science café is overturning the scientific communication, bringing science from classrooms institutional life in everyday life. The aim is not only to inform people about issues of science and technology, but, above all, to teach people to reason according to the scientific method, making them active participants in the discussion and thus creating a collective intelligence that serves not only the public, but also to experts, that learn how to listen to the public and how to talk with it. Therefore, the concept of science café has more to do with participation rather than just with science.

The (ideal) working of a science café is summarized in Fig. 1, that exploits the double meaning (in Italian) of the word "caffè": it means both a café (place) and coffee, so a *caffè-scienza* may stand for a discussion or a special coffee brand.

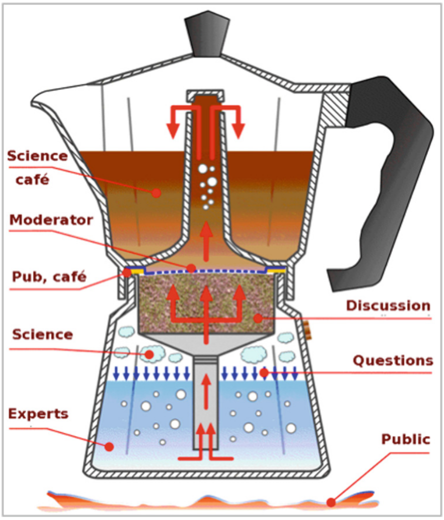

Fig. 1. An illustration of the ideal workflow of a science café: the real "motor" is the public that ask questions to the experts and though a discussion moderated and stimulated by an animator is able to distil the collective intelligence, here represented by a good Italian coffee. Image by the authors.

Science cafés quickly blossomed and nowadays they are present in almost every part of the world [19–25].

In Italy, the first science café wad held in Florence in 2004, organized by our just-funded association *Caffè-Scienza* [26]. Up to now the *Caffè-Scienza* association organized about 200 events, in many different locations. We also launched an Italian science café network in which the participants can exchange information, event calendar, videos and so on [27].

3 The Italian Public of Science Cafés

In 2014, for a project funded by the Italian Ministry of Education [28], we commissioned (and partially conducted) an investigation [29] about the public of four popular science cafés in Italy: Florence, Prato, Rome and Bari (although the science cafés in Florence and Prato are run by the same association, their public is completely different).

The survey was filled by 230 participants, of which 112 men and 108 women. 10 people did not indicate their gender. The age of participants ranges almost evenly (40–50 responses) from 20 to 60 years old, with about 8 people below 20, 12 in the range 60–70 and 10 above 80, without any significant differences between men and women.

The education level was relatively high: The vast majority owned at least a high school degree, and more than half (93) had a university degree or a post-graduate degree (71). Among women the percentage of those with a postgraduate title is greater (42 vs 29). The 10 people who did not enter their genre not even entered their qualification.

The public in Florence and Rome are more similar, while the public in Bari and Prato are younger. In particular, the audience of Prato consists of a large number of school students.

The principal channels for becoming aware of the existence of a science café are the world of mouth (132), followed by poster and flyers (24) and the web (19).

The vast majority (186) declared to prefer live science cafés, with only three declaring to prefer the online version, 27 both and 14 not expressing any preference. 198 people never followed an online science café, 24 once and 8 more than once. Notice that we started our experiments about the streaming of our science cafés in the year 2011, and Rome started in 2012.

In the following years we experienced a steady growth of web contacts and of views of the recorded events. The online public has increased slowly, only recently reaching a persistent presence. For what concerns the last events, the on-line participants constitute the 5–30% of the total public. As for the face-to-face public, the turnover of the online public is about 50%. Many attendees living in Florence alternate between online and face-to-face participation.

4 Towards Science Cafés 2.0

The original (British) formula of science cafés, does not make use any technological support, not even a microphone. However, this approach severely limits the number of attendees, and, therefore, the benefits (people reached) versus cost (effort spent by organizers and speakers) ratio. Indeed, when a science café just comes to life, the participation is generally limited, and it can be run *spartanly*, just gathering people in a pub. But, as our history taught, the number of attendees can rapidly increase, and therefore some technology aid becomes necessary.

We had a series of challenges in mind: how to keep growing the audience still keeping the management on a voluntary basis; how to favour the debate even for distant participants; allow people to exploit the "collective intelligence" produced during a debate also after its conclusion, experimenting with several media.

We have essentially studied three cases: the technology for a successful run in presence, after-event follow-up the involvement of people at distance, like audio and video streaming.

The main requirement for a successful face-to-face meeting is a good audio system which is also useful for picking the audio stream (possibly from a digital mixer), and a slide projector, which can be used for distant participation of speakers and for occasionally show the online interface.

For the follow-up, we started resuming the discussion publishing it on the web, then we switched to audio-only recording and finally to full video recording with editing for inserting the slides at full resolution.

In recent years, and especially within the European Projects SciCafé [30] and SciCafe2.0 [31], various technological solutions have been tried for the streaming, and after some tests we decided to allow real-time interaction via a textual chat due to the prohibitive requirements of a video feedback.

All these improvements, while requiring special skills, has allowed to broaden the audience to hundreds of users.

For what concerns video streaming, we experimented with several tools, from Livestream [32] to a custom platform developed within the SciCafé European project [30] and finally to YouTube [33]. We decided to stick with the last one (Google hangout on air) but adding a custom chat, since the YouTube stream allows participants to chat during the live event (and offers some tool for moderating the discussion), but, contrary to our chat, the discussion is not threaded and so questions and answers or comments got mixed. The chat (Fig. 2) was developed within the SciCafe2.0 project [31].

Fig. 2. The SciCafe2.0 threaded chat.

4.1 The SciCafe2.0 Platform

The European Project SciCafe2.0 had the goal of extending the method of science cafés to crowdsourcing, grass-rooting and e-democracy, and as a supporting tool for existing projects in the field of collaborative platform. So, during this project, we developed a platform partly based on the requirement coming from our science café experience [34].

The platform allows to embed the video streaming from a public Google Hangout On Air (i.e. streaming on YouTube [33]) and to add our chat in the same web page (in our case on the site of our no-profit association [26]), so that people can see the event and participate with questions (Fig. 3).

During the event, one person is appointed to follow the discussion with the camera for the streaming while another one is charged to follow the chat discussion with the distant participants (Fig. 4). It is highly important to give feedbacks to them so that they do not feel to be alone, since they cannot rely on visual clues and feedbacks.

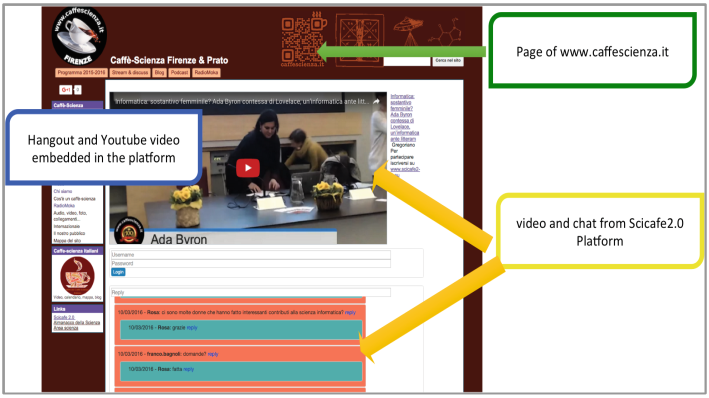

Fig. 3. Video and chat embedded in the Caffè-Scienza association web site [26].

Fig. 4. A science café in Florence: on the right one can see the technical arrangement.

The most important task of the moderator is to gather questions from distant public and report them to the expert. This has also the effect of making all public and experts aware of being followed by people on the Internet.

We tested this modality more than 40 times. However, for the last two events of the last season [35] we had to switch to the standard YouTube chat due to a software problem of our platform, and we experienced a sudden amount of participants, with a rich side-discussion on the chat.

Indeed, people who are notified of the streaming directly from YouTube, not coming from our web site, may be induced to participate seeing other people actively discussing. We experienced for the first time also the problem of moderating people participation online, some of whom had to be suspended or banned.

Let us analyse one of these events. Despite or thanks to the technicality of the topic *The return of the two-strokes engine* (*Il ritorno del motore a due tempi*) [36], we experienced a very high number of people in presence and about 30 followers on line with a maximum of 35; many compared to the average of 5–10 people of past events.

During the evening, however, some online people began to degenerate using scurrilous and unapproved language, which had never happened in the chat platform. The moderator had to temporary exclude, and then completely ban these people.

Fig. 5. YouTube page with the list of videos of our last events and the number of visualizations.

The chat on the SciCafe2.0 platform do not have a tool to ban people "non-politically correct" but to access to the platform one needs to register and perhaps this requirement prevented inappropriate behaviours. Differently from our platform, the live YouTube chat disappears after the event, preventing further analysis. After one month, the video of the two-strokes engine is still one of the more viewed, as shown in Fig. 5.

The repository on YouTube channel allows the asynchronous vision of the event, and we noticed that there are many views immediately after the ending of an event, possibly people who joined after the beginning and decided to start over or who preferred to skip parts of the conference and jump to the most interesting portion: the average duration of continuous views is about 6–10 min (with a large dispersion) over a total length of a couple of hours. Another interesting fact is that almost every video continues to be viewed after the end of the streaming, even for years (Fig. 6).

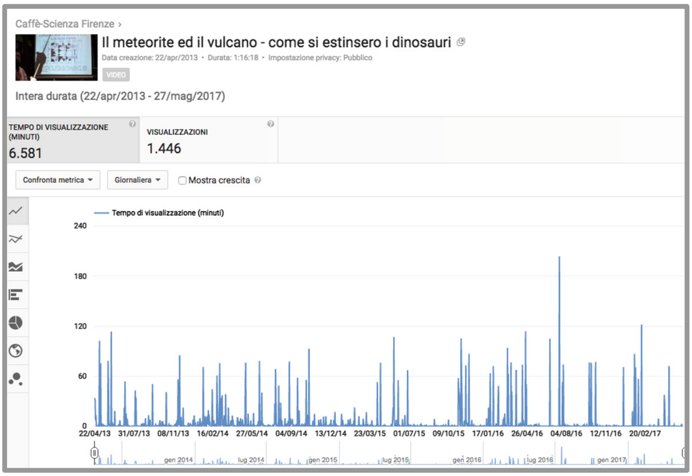

Fig. 6. A time plot of visualizations of one video. One can notice that the video continues to be viewed after years, without any evident fading out.

The SciCafe2.0 platform also allows the archiving of fluxes of information (video, documents, links) with attached discussions (Fig. 7), an option useful for managing projects.

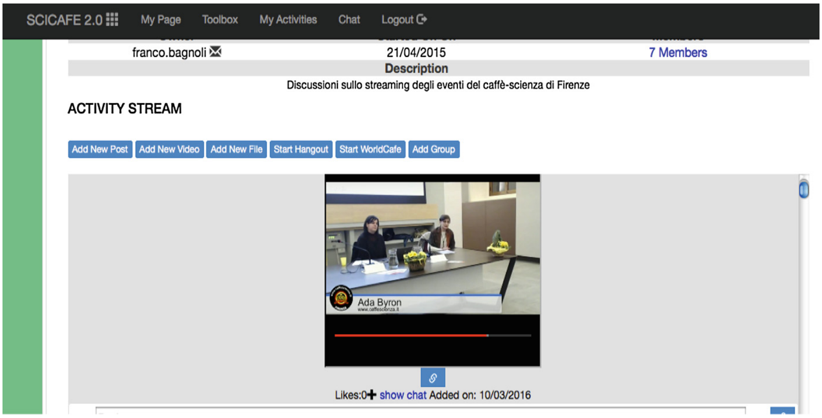

Fig. 7. An event viewed from the SciCafe2.0 platform.

4.2 Application of the Science Café Methodology to Other Communication Events

We applied the lessons learned with "traditional" science café events (i.e., prompting for questions by the public as soon as possible) to other typology of *scenarios*.

For example, we noticed that participants hesitate in asking questions about some technical topics, albeit being attracted by them, due to the fact that the attendees had little previous knowledge of the subject. We therefore decided to adapt the science café formula to allow a longer introduction (in general with just one speaker), delaying the questions. We call this modality *cafferenza*, a mix between *caffè-scienza* and *conferenza* (conference).

We used the streaming by our platform also during workshops and conferences to allow distance participation of the speakers and/or of the public.

Another interesting experiment in which we use streaming is the dissemination of the results of European projects: some events have been used to promote the public engagement and the participation in the project.

5 Lesson Learned and Conclusions

In this paper we have described the application of some techniques used in the Web 2.0 to the science café, a methodology of participative science communication.

We experimented the use of web technologies and social networks, the simultaneous use of various communication channels (mash-up of web, videos and chat) and the podcast of the material produced. To our surprise, the podcasts of the past videos are still viewed after years, without any evident fade out of the audience.

Although the degree of innovation of our experiments is surely limited from a technological point of view, there are few other science cafés that performed the streaming of their events (as far as we know, Alessandria, Livorno and Rome), all using our platform.

We adapted the science café methodology to different scenarios, since we realized that there is not a unique formula for all subjects and public.

We experienced a sudden amount of online participants when we switched to the YouTube chat due to a temporary problem in our platform, with moderation problems. We are actively considering mirroring the two chats in order to provide both visibility and threaded discussions.

In conclusion we can say that the online experimentation had a good impact on the public of the science cafés in Florence: using the web, we have witnessed an increase in both the number of users of their participation in the dialogue, not only online but also in presence.

Acknowledgements. The support of the European Commission under the FP7 Programme *Collective-Awareness Platforms*, Grant Agreement no ICT-611299 (SciCafe2.0) and of MIUR L. 6/2000 "la scienza ha fatto rete" is gratefully acknowledged.

References

1. Kennefick, D.: Einstein versus the physical review. Phys. Today **58**(9), 43 (2005). doi:10.1063/1.2117822
2. Royal Society: Report on the Public Understanding of Science (PUS) (1985). https://royalsociety.org/topics-policy/publications/1985/public-understanding-science/
3. Bodmer, W.: The Public Understanding of Science, Report-Council of Royal Society (1985)
4. Bodmer, W.: Public understanding of sciences: the BA, the Royal Society and COPUS. Notes Rec. R. Soc. **64**, S151–S161 (2010). doi:10.1098/rsnr.2010.0035
5. Rutherford, J.: Windows on the world of science: preparation and opportunity. Presented at the International Conference Science before the Public. Humanistic Culture and Scientific Technological Development, Salamanca, Spain (2002)
6. Burns, T.W., O'Connor, D.J., Stocklmayer, S.M.: Science communication: a contemporary definition. Public Underst. Sci. **12**(2), 183–202 (2003). doi:10.1177/09636625030122004
7. Carrada, G.: Comunicare la Scienza. Kit di sopravvivenza per ricercatori. (Sironi Editore 2005)
8. Hilgartner, S.: The dominant view of popularization: conceptual problems, political uses. Soc. Stud. Sci. **20**, 519–539 (1990). doi:10.1177/030631290020003006
9. Trench, B.: Science communication and citizen science: how dead is the deficit model? In: Proceedings of the Public Communication of Science and Technology Conference 2006 - PCST9, COEX, Seul, Korea, 17–20 May 2006. https://pcst.co/archive/pdf/Trench_PCST2006.pdf
10. Office of Science and Technology and the Wellcome Trust: Science and the public – a review of science communication and public attitudes to science in Britain. Public Underst. Sci. **10**(3), 315–330 (2000). doi:10.3109/a036873
11. Bauer, M.W.: The evolution of public understanding of science—discourse and comparative evidence. Sci. Technol. Soc. **14**, 221–240 (2009). doi:10.1177/097172180901400202

12. Worcerster, R.: Science And Society: What Scientists And The Public Can Learn From Each Other. MORI report (2001)
13. Pollock, J., Steven, D.: Now for the science bit – concentrate! – a report into the public understanding of science. River Path Associates (1997)
14. Miller, S.: Public understanding of science at the crossroads. Public Underst. Sci. **10**, 115–120 (2001). doi:10.3109/a036859
15. Hughes, J.: Science and society: a public role for the history of science in the 2000 s? Br. Soc. Hist. Sci. Newslett. **62**, 20 (2000)
16. House of Lords, Select Committee on Science and Technology, Third report: Science and Society, 23 February 2002
17. Research Councils and Office of Science and Technology UK. Dialogue with the public: practical guidelines (2002)
18. Dallas, D.: The café scientifique. Nature **399**, 120 (1999). doi:10.1038/20118
19. Ferris, D.: Science Cafés tap nation's fascination with research and discoveries. Wired (2007). https://www.wired.com/2007/12/science-cafs-tap-nations-fascination-with-research-and-discoveries/
20. Mutheu, J., Wanjala, R.: The public, parasites and coffee: the Kenyan Science Café concept. Trends Parasitol. **25**(6), 245 (2009). doi:10.1016/j.pt.2009.03.002
21. Sink, M.: Science Comes to the Masses (You Want Fries With That?), The New Yourk Times, 21 February 2006. http://www.nytimes.com/2006/02/21/science/science-comes-to-the-masses-you-want-fries-with-that.html
22. Chen, D.-S., Wu, C.-L.: Introduction: public participation in science and technology in east Asia. East Asian Sci. Technol. Soc.: Int. J. **1**, 15–18 (2007). doi:10.1007/s12280-007-9005-6
23. Bagnoli, F., Dallas, D., Pacini, G.: Sipping Science with a Science Café, Lulu.com (2012). https://sites.google.com/site/scicafewebbook/, ISBN 9781105846908
24. Nielsen, K.H., Balling, G., Hope, T., Nakamura, M.: Sipping science: the interpretative flexibility of science cafés in denmark and Japan. East Asian Sci. Technol. Soc.: Int. J. **9**, 1–21 (2015). doi:10.1215/18752160-2832109
25. Bagnoli, F., Pacini, G.: Sipping science in a Café. Procedia Comput. Sci. **7**, 194–196 (2011). doi:10.1016/j.procs.2011.09.070
26. Caffè-Scienza Association in Florence. http://www.caffescienza.it
27. Italian network of science cafés. http://www.caffescientifici.it
28. Project MIUR ex 6/2000 "la scienza ha fatto rete"
29. Castellani, T., Valente, A: Rapporto sul pubblico dei caffè scienza italiani (Report on the Public of the Italian Science Cafés) (2015). http://www.irpps.cnr.it/e-pub/ojs/index.php/wp/article/view/74
30. SciCafé project web site. http://www.scicafe.eu
31. SciCafe2.0 web site. http://www.scicafe2-0.eu
32. Livestream web site. https://livestream.com/
33. YouTube channel of the Caffè-Scienza Association. https://www.youtube.com/user/caffescienza
34. Badii, A., et al.: Collective awareness platforms and digital social innovation mediating consensus seeking in problem situations. In: Bagnoli, F., Satsiou, A., Stavrakakis, I., Nesi, P., Pacini, G., Welp, Y., Tiropanis, T., DiFranzo, D. (eds.) INSCI 2016. LNCS, vol. 9934, pp. 55–65. Springer, Cham (2016). doi:10.1007/978-3-319-45982-0_5
35. Caffè-Scienza 2016–2017 program. http://www.caffescienza.it/programma-2016-2017
36. Florence Science Café: The return of the two-stokes engine, 4 May 2017. http://www.caffescienza.it/programma-2016-2017/motore-due-tempi

Online Policy, Politics and Co-creation

Conceptualizing ICT-Enabled Co-creation of Public Value

Aelita Skaržauskienė and Monika Mačiulienė(✉)

Mykolas Romeris University, Ateities 20, Vilnius, Lithuania
{aelita,maciuliene}@mruni.eu

Abstract. Traditional views on the public value creation focus on the public sector organizations as sole initiators of the value creation processes. The rise of interactive Information and Communication Technologies (ICT), however, opens new opportunities for broader engagement of civic stakeholders in the public value generation. The concept of co-creation is seen as a new framework describing the shift from considering organizations as the definers of the value to a more inclusive and collaborative processes involving the end-users and other external actors. The article proposes a conceptual framework providing holistic integration of current research efforts on the co-creation of public value by focusing on the initiatives originating outside governmental entities. The conceptual framework provides understanding on how ICT-enabled co-creation should be utilized in the generation of the public value.

Keywords: Public value · Co-creation · ICT

1 Introduction

Traditional views on the public value creation focus on the public sector organizations as sole initiators of the value creation processes. The rise of interactive Information and Communication Technologies (ICT), however, opens new opportunities for broader engagement of civic stakeholders in public value generation. The concept of co-creation is seen as a new framework describing the shift from considering organizations as the definers of value to a more inclusive and collaborative processes involving end-users and other external actors. While the concept originated in the business context, the co-creative strategies could benefit the public and the civic sectors too. The contributions by the citizens, private and non-governmental entities are increasingly considered as an effective tool for tackling multifaceted issues of contemporary society in terms of co-designing solutions to the problems together with the experts and public officials. However, there is a lack of clarity in the literature regarding the forms co-creation can take in the public sector and the research surrounding it. The article proposes a conceptual framework providing holistic integration of current research efforts on co-creation of public value by focusing on initiatives originating outside governmental entities. The conceptual framework provides understanding on how ICT-enabled co-creation should be used in developing public value. The study contributes to the initial phase of the C^3PLACES research project aiming at exploring co-creative strategies and tools increasing the quality of public open spaces.

I. Kompatsiaris et al. (Eds.): INSCI 2017, LNCS 10673, pp. 93–100, 2017.
https://doi.org/10.1007/978-3-319-70284-1_8

2 The Theoretical Premises of the Framework

The conceptualization of ICT-enabled co-creation is based on the combinations of the principles of ICT-enabled public sector theory, governance theory and co-creation theory (See Fig. 1). These research fields have a joint emphasis on the collaborative processes enabled by technological advancements, but they differ in focus. First, the contributions of co-creation theory are discussed. Authors have highlighted the interactive and networked nature of value creation both in the business [1, 2] and the public sectors [3, 4]. In line with this trend, the co-creation theory of Service science emerges forming the base for understanding the relationships between different actors [5, 6]. The motivation to create partnerships comes from the recognition that organizations can accomplish, what each partner cannot accomplish alone, by maximizing the influence, creating collective resources, and removing duplication of the efforts. Although the public sector can in principle create the public value on its own, its potential to do so is greatly enhanced and extended by direct cooperation with other actors, or by facilitating public value creation by other actors on their own. Service science suggests that value emerges when a number of entities work collectively to create mutual benefit by granting access to one another's resources including people, technologies, organizations and information.

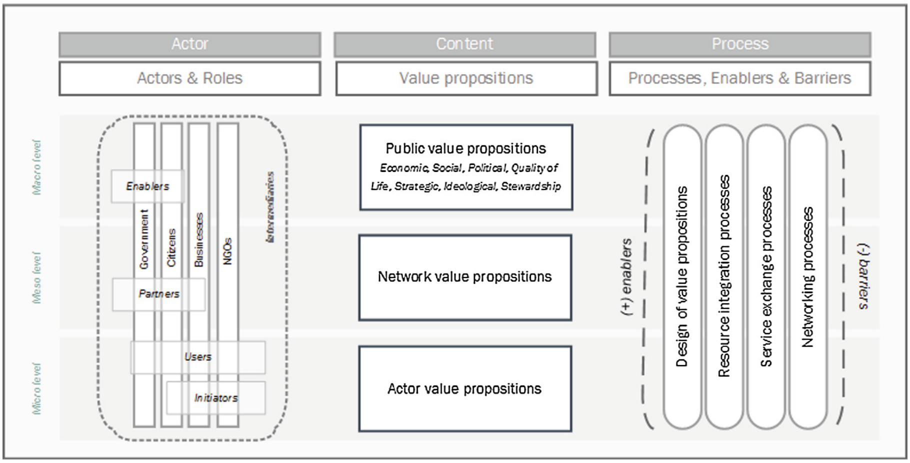

Fig. 1. Conceptual co-creative ecosystem framework (developed by authors, 2017)

ICT-enabled public sector theory contributes to the conceptualization through Government 2.0 approach and provides the basis for understanding value propositions the governments can provide to the civic society in terms of open data and facilitation of transparency and openness. Government 2.0 is a novel way to define the use of Web 2.0 technologies in socializing the government services, processes and data. At the center of this approach is the understanding of the government as a platform providing data and support technology to increase the participation and transparency [7].

The governance theory contributes to the conceptualization through the New Public Governance and public value approaches. New Public Governance approach explains the context and the need for changing the power balance and enabling collaborative practices in the creation of public value. It is characterized by the network perspective which takes into account the inputs by non-governmental organizations, private sector, educational organizations, international institutions in the processes of governance. In NPG perspective, traditional hierarchies are replaced by new organizational forms and involves sharing the infrastructures, processes, data, assets, knowledge, resources, content and tools [8]. The notion of public value is of a key importance in the deliberations on NPG. The concept was coined by [9] who suggested that "the value of public services is not limited to the quality or efficiency of those services, but also pertains to the actual social and economic improvements they create for the public". Based on the rational put forward by [9], the value of public services is not limited with the efficiency and quality but also deals with the social and economic improvements they create for the society. Following this logic, public value is defined as the total societal value that cannot be monopolized by individuals, but is shared by all actors in the society and is the outcome of all resource-allocation decisions.

Taking into account the discussion on theoretical developments in the fields of Co-Creation, ICT and Governance, ICT-enabled co-creation of public value is defined as a system driven by the goal of generating public value through the use of ICT and co-creation between the government sector, the private sector, and the civil society.

3 The Forms of ICT-Enabled Co-creation

The governments are becoming more user-centric which is expressed through the focus on servitization of the sector [10, 11], design thinking [12–14], online public engagement initiatives [15–17], etc. In addition, there has been an increase in digital solutions oriented towards creation of public value developed by entities – civil society organizations, individual citizens and businesses – outside government. Two general directions of ICT-enabled co-creation of public value research can be identified – top-down and bottom-up – differing on the understanding of the roles of governmental entities. Current research efforts bundle these different approaches to co-creation together and structured distinction is needed. Mostly because, government-led platforms "can be vulnerable to institutional biases and rationale, and the resulting tools may be built with inherent assumptions concerning the users' needs" [18, p. 14].

Top-down co-creation approach refers to the implementation, design, and evaluation of public services, participation in government-initiated platforms, data and content contribution, improvement of existing processes and services, user-centric approaches to service design. The literature provides variety of taxonomies for understanding government-led co-creation [19–22]. Approach by [23] provides the broadest classification and determines three types of co-delivery initiatives: co-design (allows citizens to participate in development of new service or policy), co-production (involves citizens in creating a service), and co-delivery (involved citizens in delivering the services). The literature on co-creation of public value mostly focuses on the top-down approach by analyzing the processes in the governmental structures and suggest managerial approaches for government officials on how to create more open governance systems and how to involve citizens, businesses, and non-governmental organizations into these systems.

The bottom-up co-creation approach refers to the platforms emerging from outside the governmental sector. Such technologies are not necessarily created with the aim of being disruptive, but are designed by and for the citizens, using open data, collaboration technologies in innovative ways that can improve the existing channels of information and communication previously controlled by institutions alone [24]. Based on [25], it refers to the innovation outside the government control and may have significant impact on the communications between civic society and governmental entities. It is driven by the increased innovative use of ICTs and changed expectations of citizens towards the interactivity, simplicity, comfort and integration of the public services. Governments are struggling to match the pace of private sector in creating improved services and value propositions. Hence, growing number of platforms and applications lie at the intersection of technology and governance. Citizens, NGOs and private sector cannot re-create the services offered by the governmental organizations due to a number of reasons – the resources, the scale or security issues. Only governmental organizations can create large-scale projects such as eHealth or eVoting. However, the civil society can contribute in creating smaller tools increasing the transparency and accountability or building the communities. The research on bottom-up co-creation is limited and remains at initial phase. The field lacks generalization, established theoretical models and empirical evidence. While the literature [26–29] within this stream provides multiple examples of civic society initiatives that have applied principles of co-creation, there exists limited amount of studies regarding certain activities that should be undertaken in order to enable co-creative processes.

4 Conceptualizing Bottom-Up Co-creation of Public Value Enabled by ICT: Building Block of the Framework

The framework (Fig. 1) we are proposing provides holistic view into the public value co-creation. The perspective of Service Science provides a framework for describing complex relationships between the public entities, private entities and civil society. [30] view service as the fundamental unit of exchange and explain it as the application of competences by one party for the benefit of another. The Service science suggests that value emerges when a number of entities work collectively to create mutual benefit by granting access to one another's resources including people, technologies, organizations and information. Interacting entities form service ecosystem consisting of several or many service systems connected by a network. The metaphor of an ecosystem arrived from biology and refers to a system of interacting entities which depend on one another for effectiveness and survival [9]. In the context of this research project, service ecosystem refers to a system in which actors work together to achieve mutual benefit – public value.

The actors in the proposed framework can create value on three levels – micro, meso and macro. Micro level refers to the direct service-for-service exchange with the end-users. Meso level refers to indirect service exchanges with other stakeholders (e.g.

partners, competitors). Macro level refers to the complex relationships between different systems with diverse interests co-creating public. The exchanges between the actors in different levels of are needed because no one actor has all the resources needed to reach identified goals [31]. In order to understand how public value is created on micro, meso and macro levels, three dimensions – actors, content, processes – were developed allowing to categorization of the entities involved and ways they co-create public value.

The processes dimension includes deliberations on the patterns of design, management and collaboration in co-creating public value. Service science provides a lens to view actors in a system of other actors co-creating value through resource integration and service provision [32]. A service ecosystem has structural integrity because each entity (economic and social actors) has competencies (used to offer and provide service to others), relationships (with other actors) and information that is shared. Value propositions are then used to connect one actor with other interested actors within that service ecosystem. To this end, the actor that develops the most compelling value proposition (offers a connection between competences and relationships) will perform the best; however, this relative performance advantage will be fleeting unless the organization learns to revise its value propositions in response to changing customer, supplier, and other stakeholder requirements. Research [25, 33–40] suggests that the co-creation of public value is influenced by a variety of preconditions (enablers and barriers) on the ecosystem levels – macro (strategic policy framework, infrastructure for openness, transparency and accountability, features of civic society, institutional support, open attitude of public official, lacking), meso (interoperability between governmental entities, heterogeneity of actors, embeddedness in networks, offline engagement strategies), micro (integration of external input, risk aversion of actors, clear incentives). However, most of the research efforts focus on the preconditions on the macro level and more theoretical and empirical investigations are needed on meso and micro levels.

The actors dimension refers to the entities (i.e. individuals and organizations) participating in the service ecosystem, their roles and resources. [41, p. 285] suggests that "although collaboration has the potential to produce powerful results, not all collaborations realize this potential. Many collaborations fail to produce innovative solutions or balance stakeholder concerns, and some even fail to generate any collective action whatsoever". Hence, the understanding of the actors involved in co-creation and the roles they can perceive is crucial. The concept of roles allows to get insights on the ways actors collaborate in service systems. Despite the diversity of actors involved in any ecosystem, it is possible to identify different types of actors, segment them and understand the nature of their relationships within defined context. Four types of actors (government entities, citizens, business entities, NGOs) and five roles they can perceive have been identified within research literature. The actors can be initiators [24, 25, 27, 45], users [25, 38, 42], partners [25, 45–47], enablers [45] and intermediaries [43, 44].

The content dimension includes deliberations on the goals and objectives of the actors involved. Knowing why individuals and organizations build platforms, and why citizens participate in them, can guide the organizations and civic leaders in fostering ICT-enabled platforms. Central concept of this dimension is value proposition. To realize the value proposition, a firm must co-create value with other actors in the system by direct interactions benefiting both sides. Value propositions indicate how the actors involved could co-create value by integrating ecosystems their resources because the actor cannot deliver the value, but only offer value propositions. Value proposition refers to a promise of value to be delivered, communicated, and acknowledged. The typology of value proposition is based on the works of Moore's Public Value theory and elaborations of this theory by [43, 44]. A well-articulated value proposition can help streamline organizational decision-making. [43] suggests seven types of public value: Economic, Social, Political, Quality of life, Strategic, Ideological, and Stewardship. [44] added deeper perspective to understanding public value creation in ecosystems and suggested that the organizational strategy to public value involves several layers. This notion coincides with the micro, meso and macro approach. Hence, the public value is co-created on the macro level since it includes deliberations on the context and larger social constructs (government structures, civic society, etc.). Meso level provides insights on the stakeholder network benefits. Micro level deals with value offerings for the individual actors. By distributing value propositions though three levels, the framework allows to understand the value of ICT-enabled co-creation for people, organizations and society.

5 Recommendations for Further Research

The conceptual framework provides a holistic view into the co-creation processes in the public value generation. The conceptual framework presented is designed based on the theories and past investigation on the subject. It provides an interpretative approach to the social reality. Hence, more empirical investigations are needed to test its consistency with the reality aimed at testing its consistency with reality. Empirical investigations are needed to gain more insights on the interrelations between the elements and the validity of the proposed framework. Three empirical research directions can be identified: Analysis of the processes in the ecosystem (What influences the processes of resource integration, service provision and development of sustainable value propositions in the ecosystem? What factors enable these processes?), the actors in the ecosystem (What are the actors involved in the processes of ecosystem and what are they roles? Do the actors identified in the literature matched the reality of civic technologies? How the actors in micro, meso and macro levels connect?) and the content of the public value co-creation (What public value civic technologies offer to society? What value propositions gain most value in ecosystems?).

Acknowledgements. This research (C^3PLACES - using ICT for Co-creation of Inclusive Public Places) was funded by a grant (No. 7135959) from the Research Council of Lithuania.

References

1. Akaka, M.A., Vargo, S.L., Lusch, R.F.: The complexity of context: a service ecosystems approach for international marketing. J. Int. Mark. Am. Mark. Assoc. **21**(4), 1–20 (2013)

2. Gronroos, C.: Value co-creation in service logic: a critical analysis. Mark. Theory **11**, 279–301 (2011)
3. Hänninen, H.: Management of Value Co-creation in Public Service Networks - Case City of Helsinki, p. 99 (2012)
4. Wise, S., Paton, R.A., Gegenhuber, T.: Value co-creation through collective intelligence in the public sector: a review of US and European initiatives. Vine **42**(2), 251–276 (2012)
5. Vargo, S.L., Lusch, R.F., Akaka, M.A.: Advancing service science with service-dominant logic: clarifications and conceptual development stephen. In: Maglio, P., Kieliszewski, C., Spohrer, J. (eds.) Handbook of Service Science. Service Science: Research and Innovations in the Service Economy, pp. 133–156. Springer, Boston (2010). doi:10.1007/978-1-4419-1628-0_8
6. Vargo, S.L., Akaka, M.A.: Service-dominant logic as a foundation for service science: clarifications. Serv. Sci. **1**(1), 32–41 (2009)
7. O'Reilly, T.: Government as a platform. Innovations **6**(1), 13–40 (2011)
8. European Commission: Consultation on directions for ICT-driven public sector innovation at European Union level Research and innovation in Horizon 2020 (2013)
9. Moore, M.H.: Creating Public Value: Strategic Management in Government, p. 402. Harvard University Press, Cambridge (1995)
10. Denhardt, J.V., Denhardt, R.B.: The new public service revisited. Public Adm. Rev. **75** (October), 664–672 (2015)
11. Osborne, S.P., Radnor, Z., Nasi, G.: A new theory for public service management? Toward a (public) service-dominant approach. Am. Rev. Public Adm. **43**(2), 135–158 (2012)
12. Leavy, B.: Collaborative innovation as the new imperative–design thinking, value co-creation and the power of 'pull'. Strateg. Leadersh. **40**(2), 25–34 (2012)
13. Allio, L.: Design Thinking for Public Service Excellence (2014)
14. Brown, B.T., Wyatt, J.: Design thinking for social innovation. Stanford Soc. Innov. Rev. **Winter**(Winter), 30–35 (2010)
15. Gatautis, R.: Creating public value through eparticipation: wave project. Ekon. ir Vadyb **15**, 483–490 (2010)
16. Peixoto, T., Fall, M., Sjoberg, F.: Evaluating Digital Citizen Engagement: A Practical Guide. The World Bank, Washington, D.C (2016)
17. Giest, S., Koene, A., Vallejos, E.P., Pitkänen, O., Fosci, M.: Online spaces for urban citizen engagement: a comparison of civic apps. In: Data for Policy Conference, pp. 1–6 (2016)
18. Rumbul, R.: Developing transparency through digital means? Examining institutional responses to civic technology in Latin America. JeDEM **8**(3), 12–31 (2016)
19. Magno, F., Cassia, F.: "Public administrators' engagement in services co-creation: factors that foster and hinder organisational learning about citizens. Total Qual. Manag. Bus. Excell, pp. 1–12. January 2014 2015
20. Sherriff, L.: Sensing the future: how will smart city principles and technology enable citizen co-creation in public policy-making, consent processing and service provision? In: 2015 Conference on Back to the Future, New Zealand Planning Institute, pp. 1–15 (2015)
21. Hilgers, D., Ihl, C.: Citizensourcing: applying the concept of open innovation to the public sector. Int. J. Public Particip. **4**(1), 67–88 (2010)
22. Cabrera, A.B.: New Public Governance, Co-production & Third Sector Social Services, pp. 1–23. October 2010
23. Kannan, P.K., Chang, A.-M.: Beyond citizen engagement: involving the public in co-delivering government services. Collab. Across Boundaries IBM Cent. Bus. Gov. Washingtn, D.C (2013)
24. Suri, M.V.: From crowdsourcing potholes to community policing: applying interoperability theory to analyze the expansion of 'Open311'. Berkman Cent. Res. Publ. No. 2013–2018 **7641**, 1–19 (2013)

25. Rumbul, R.: ICT and citizen efficacy: the role of civic technology in facilitating government accountability and citizen confidence. IFIP Adv. Inf. Commun. Technol. **481**, 213–222 (2016)
26. AVINA Foundation: Civic technologies as tools for change: is balance between online and offline actions a necessary condition for achieving social change through civic innovation? (2015)
27. Knight Foundation: The Emergence of Civic Tech: Investments in a Growing Field, p. 30. December 2013
28. Baeck, P., Bria, F.: Digital Social Innovation: What It Is And What We Are Doing, London (2014)
29. Marian, D., Monteyne, D.: Urban co-creation: envisioning new digital tools for activism and experimentation in the city. Chi **2011**, 1–4 (2011)
30. Vargo, S.L., Maglio, P.P., Akaka, M.A.: On value and value co-creation: a service systems and service logic perspective. Eur. Manag. J. **26**, 145–152 (2008)
31. Frow, P., Payne, A.: A stakeholder perspective of the value proposition concept. Eur. J. Mark. **45**(1/2), 223–240 (2011)
32. Lusch, R.F., Vargo, S.L.: Service-dominant logic: a necessary step. Eur. J. Mark. **45**(7/8), 1298–1309 (2011)
33. Kearns, I.: Public Value and E-Government. Institute for Public Policy Research, London (2004)
34. McGee, R., Gaventa, J., Barrett, G., Calland, R., Carlitz, R., Joshi, A., Acosta, A.M.: Review of Impact and Effectiveness of Transparency and Accountability Initiatives: Synthesis Report (2010)
35. Akaka, M.A., Chandler, J.D.: Roles as resources: a social roles perspective of change in value networks. Mark. Theory **11**(3), 243–260 (2011)
36. Sekulova, F.: WP3: Analysis of success factors Report on qualitative success factors (2016)
37. Capgemini, Sogeti: Digitizing Public Services in Europe: Putting Ambition into Action (2011)
38. Åkesson, M.: Role constellations in value co-creation: a study of resource integration in an e-government context (2011)
39. Voorberg, W., Bekkers, V., Tummers, L.: Co-creation in social innovation: comparative case-study on the influential factors and outcomes of co-creation. Lipse **320090**(320090), 1–31 (2014)
40. Magno, F., Cassia, F.: Public administrators' engagement in services co-creation: factors that foster and hinder organisational learning about citizens. Total Qual. Manag. Bus. Excell. **26** (11–12), 1161–1172 (2014)
41. Cobo, C.: Networks for citizen consultation and citizen sourcing of expertise. Contemp. Soc. Sci. **7**(3), 283–304 (2012)
42. Badger, E.: The next big start-up wave: civic technology (2012). http://www.citylab.com/tech/2012/06/next-big-start-wave-civic-technology/2265/. [Accessed: 14 May 2017]
43. Cook, M.: Delivering public value through transparency. In: Burke, W.M., Teller, M. (eds.) A Guide to Owning Transparency: How Federal Agencies Can Implement and Benefit from Transparency, pp. 16–21. Open Forum Foundation (2011)
44. den Ouden, E.: Innovation Design: Creating Value for People, Organizations and Society. Springer Science & Business Media, London (2011)
45. Clarke, R.Y.: Civic Tech Fuels U.S. State and Local Government Transformation (2014). http://www.accela.com/images/resources/whitepaper/idc-civic-tech-report.pdf
46. Sifry, M. L.: Civic Tech and Engagement: In Search of a Common Language (2014). http://techpresident.com/news/25261/civic-tech-and-engagement-search-common-language. Accessed 7 June 2016
47. Nambisan, S., Nambisan, P.: Engaging Citizens in Co-Creation in Public Services - Lessons Learned and Best Practices. Collaboration Across Boundaries Series (2013)

Designing a Digital Social Innovation Platform: From Case Studies to Concepts

Ines Dinant[1], Jacqueline Floch[2], Thomas Vilarinho[2], and Manuel Oliveira[2(✉)]

[1] Farapi, Gipuzkoa, Spain
ines@farapi.com
[2] SINTEF, Trondheim, Norway
{jacqueline.floch, thomas.vilarinho, manuel.oliveira}@sintef.no

Abstract. Governments in the western countries are faced with a number of growing social challenges, such as unemployment, migration, ageing population, explosion of chronic disease. Although they offer a wide range of public social services, we cannot assume that the economy will grow at a rate that can fund expanding needs for services risen by these challenges. We have to find new ways to adapt service provision and prevent social exclusion. Social innovations are new approaches to addressing social needs through engaging beneficiaries and supporting actors in the development of solutions. There is great potential in exploiting digital networks for social innovation. Supporting virtual communities and new forms of collaboration, digital networks make it possible to co-create knowledge and solutions at a wide scale. Various digital social innovation platforms have emerged in the recent years. However we observe that these platforms focus on specific areas, such as open democracy, collaborative consumption or environment, rather than providing support for a wide range of social challenges. We propose to develop a digital social innovation platform that facilitates citizens and organisations to collaboratively develop innovative social solutions. From the analysis of the current innovation processes and the expectations of two distinct cases, Cibervoluntarios (CIB) and Experts-In-Teamwork (EiT), we derive an initial set of concepts that serve as a basis for the development of a methodology and platform for social innovation.

Keywords: Social innovation · Co-creation · Open platforms

1 Introduction

A social innovation is a novel solution to a social problem. The TEPSIE project defines social innovations as innovations that are "social in their means and in their needs" [1]. Beyond solving social needs, social innovations engage and mobilise beneficiaries in the development of solutions. Social innovation has increasingly gained focus in Europe, as evidenced by the establishment of the research area CAPS (Collective Awareness Platforms for Sustainability and Social Innovation) in the European work programme FP7.

I. Kompatsiaris et al. (Eds.): INSCI 2017, LNCS 10673, pp. 101–118, 2017.
https://doi.org/10.1007/978-3-319-70284-1_9

Europe faces a number of growing social challenges, such as unemployment, migration, ageing population, explosion of chronic disease. Traditionally, the governments have provided a wide range of public social services in order to support social needs. We cannot assume that the economy will grow at a rate that make it possible to rely on public services to support the needs. How can we provide care to an increasing elderly population? How can we tackle the massive arrival of refugees? Social innovation is about engaging citizen in taking responsibility in solving societal problems. However, social innovation does not imply that citizens should solve problems alone. Public institutions and other organisations may drive social innovation with the citizen playing a major role, thus releasing institutions from carrying all responsibility.

Success stories of social innovation often illustrate initiatives taking place in local communities [2], with limited impact beyond the socio-spatial location where the innovation takes place. However, according to NESTA, things are moving slightly towards scalable social innovation, with concrete examples providing a better understanding on how innovation can scale [3]. Digital networks provide new opportunities for social innovation. Supporting virtual communities and new forms of collaboration, digital networks make it possible to co-create knowledge and solutions at a wide scale. Exploiting network effects, they make it possible to mobilize and engage people, and also to spread solutions widely. Despite the potential of digital technologies in boosting networks effects, we still lack a digital meeting place where organisations and citizens can easily find information about a wide range of social challenges, and collaborate to solve these challenges.

Differently, the H2020 SOCRATIC project [4] aims at developing a platform to support both citizens and organisations to collaboratively identify specific innovative solutions for achieving the desired social goals. The platform will allow individuals, collectives, institutions, companies or administrations to propose new challenges oriented to solve specific sustainability issues, to invite individuals or organizations to participate with innovative ideas that solve these issues and to collaboratively select and implement the most promising ideas. In order to ensure that the developed solutions can be applied in different contexts, and cover different needs, the project involves two diverse organisations:

- CiberVoluntarios Foundation (CIB) is a non-for-profit organization created and composed by Social Entrepreneurs engaging volunteers on using information technologies for social innovation and enabling citizens' empowerment. They play an active role in societal change by developing volunteer work to promote the use and knowledge of technological tools among populations with low or no access to technology and training.
- The "Experts in Team" (EiT) program at the Norwegian University of Science and Technology is a study program which involve multi-disciplinary groups of students in the role of social innovators. The students work to develop solutions to challenges brought by external "customers", e.g. NGOs or public institutions.

CIB and EiT are two distinct social contexts, which have very different work processes that involve different types of users. While students at EiT are young and have high educational backgrounds, the volunteers at CIB belong to different age groups and have different levels of education. The engagement of stakeholders varies between both social contexts, within EIT, the students' engagement is bound by the study program, while at CIB the stakeholder engagement is completely voluntary.

Through the study of these different organisations and the elicitation of their needs, the SOCRATIC project aims at developing a platform, for supporting social innovation, that is more general than if it was otherwise developed targeting a single organisation. CIB and EiT have been involved throughout the project, from inception of the core ideas to the deployment and collation of evaluation data. They contribute to the definition of scenarios, they give feedback along the incremental development of the methodology and digital platform, and they will evaluate the pilot solutions.

This paper focuses on the understanding of the work processes of the involved organisations, the roles of the users that contribute to their work, and the identification of their needs. From this understanding, we derive the main elements of the SOCRATIC concept, which serve as a basis for an ICT-platform supported social Innovation methodology that can be applied to different organizations supporting social innovation.

The paper is structured as follows: Sect. 2 presents related work; Sect. 3 introduces the research methodology; Sect. 4 describes the case studies, i.e., the work processes at CIB and EiT; Sect. 5 introduces the SOCRATIC concept based on the case studies; Finally Sect. 6 concludes and identifies relevant further work.

2 Background

The Social Innovation Process (SIP) is described in the Open Book of Social Innovation as a process consisting of six steps [5]:

- Prompts: is a step which occurs in fact before the SI process itself. It is about identifying and better understanding the societal problem to be solved by the social innovation. It builds the knowledge for the next step, the ideation, to take place.
- Ideation: this step involves the generation and refinement of ideas to solve the societal challenge. As social innovation is about innovations which are social both in their ends and in their means, it is crucial that the idea definition process is socially inclusive.
- Prototyping: is about materializing the idea in a simple manner so that it can be done quickly and with relatively little resources but, at the same time, supporting evidence gain about the idea hypothesis.
- Sustaining: this step is about bringing the innovation to market and establishing a foundation (revenue streams, operational capacity, etc.) that can support the innovation to be sustainable.
- Scaling: once the innovation is operational and has a sustainable customer base, it is time to look into growing it towards a wider audience. Such grow, can be in terms of reaching new regions or beneficiary segments for example.
- Systemic change: corresponds to a state where the social innovation solution permeates different levels of the society and changes cultures and people's mindset.

When looking for digital social innovation platforms, we identified platforms that were either centered in a specific social domain (such as accessibility [6], carbon dioxide emissions [7], etc.) or provided support to or within a single step of the SI process (such as ideation [8, 9], funding [10, 11] or project management [12, 13]). We did not identify any platform which guides innovators across the whole SI process and rely on a sound methodological foundation. The best candidates, Openideo [14] and Quirky [15] are industry-led initiatives whose results haven't been formally assessed through research.

Indeed one main challenge is that many activities in the social innovation process happen outside of the digital world. Face-to-face meetings have a big impact in enabling mutual understanding, prototyping a product (with the exception of a digital one) requires physical and tangible craftship and different communication means have different efficiencies when communicating with different target groups. For example, one may need to have physical meetings to open a dialogue with beneficiaries user groups that have not widely adopted digital technology such as elderly or low income communities. In many cases, it is also essential to understand the physical context in which the problems to be solved occur.

Digital support for facilitating social innovation should therefore focus at digitizing tasks that can be optimized or have their impact increased by being performed with assistance of computer systems. During our research, we analyzed how the SIP of two different organizations can be facilitated using digital technology and we generalized the concepts in order to provide technological solutions that support the processes of similar social innovation facilitator organizations.

3 Research Approach

The research in SOCRATIC follows the design-science paradigm [16]. While behavioral-science approaches focus on the use and benefits of a system implemented in an organization, design-science approaches seek to create information systems to solve identified organizational problems. Design-science approaches follow a recursive process allowing a gradual understanding of the problem to be solved and the improvement of solutions. The creation and assessment of IT artifacts is central for understanding and improvement. The term IT artifact is used in a wide sense and denotes various items related to the creation of information systems, such as models, methods and software prototypes. The design-science paradigm does not impose any concrete research and evaluation method. The choice of a method depends on the nature of the problem to be solved and the type of IT-artifact being created. In SOCRATIC, we plan to develop several IT-artifacts: the SOCRATIC concept presented in this paper, the intermediate version of the methodology, the intermediate version of the platform, the final version of the methodology and the final version of the platform.

As the first step of the research, presented in this paper, we aimed at understanding the two organisations that serve as a basis for the requirements of SOCRATIC and which will pilot SOCRATIC at the end of the project. To that end, we followed an exploratory case study strategy [17]. We conducted two exploratory case studies, one in each of the pilot organisations. The purpose of these case studies was to understand the needs of the organisations and their expectations to the SOCRATIC platform. The research questions we had were:

- How do the organisations currently support the social innovation process?
- How can a digital platform facilitate and enhance the current process?

The two pilot organisations involve users with different backgrounds (e.g., age and education). They engage in innovators in different manners and focus on different sustainable development, thus allowing us to generate broader conclusions than involving a single organisation.

Further, following the collection and analysis of data, we developed a set of pilot scenarios together with key stakeholders in each pilot organisation in order to concretize the needs.

3.1 Data Collection

We collected data from the pilot organizations regarding our research questions via two main steps:

- First, we performed an analysis of the documentation available in the public domain, secondary sources, e.g. web pages of the organisations, training materials. This provided us a baseline understanding of the context and to prepare for the next step. That allowed us to identify the key actors to involve in the next step and to design guidelines for it.
- Second, in-depth interviews, primary sources, involving key stakeholders in each pilot organisation were carried out. In order to respond to the main research questions presented above, these interviews were conducted in a semi-structured way, based on a common script. The interviews questions were elaborated as to give light on the following topics related to our research questions: (1) the form of the social innovation process followed by the organisation, (2) the actors involved in the innovation process as well as the nature of their participation and the interaction between these actors, (3) the ICT support in the innovation process, (4) the current and foreseen challenges related to process, participation and digital support as well as ideas for potential solutions, and, (5) the approach for evaluating the activities carried out.

3.2 Data Analysis

The analysis of the collected information has followed the thematic analysis methodology [18]. This methodology supports working with a wide range of research questions through the collection and analysis of primary sources and secondary sources, and in a specific data-set. While primary sources refer to self-produced content, such as semi-structured interviews and observations, secondary sources refer to content produced externally, such as bibliography. A specific data-set means the specific context in which the research takes place, in which the information is going to be collected.

The themes of the analysis were based on the research questions and theirs codes were created and refined during the analysis process based on the elements detected to understand in detail the main themes. These are the variables attached to those. The list below presents the mapping between the different codes, themes and research questions:

- (RQ1) "How do the organisations currently support the social innovation process?
  - (T1) The understanding of social innovation among the different contributors in the organisation
    C1.1: how is innovation understood;
    C1.2: how is innovation transmitted,
    C1.3: existence of shared understanding
  - (T2) The organization's innovation process
    C2.1: actors (their interactions and roles in the process)
    C2.2: ICT as support to the process
- (RQ2) "How can a digital platform facilitate and enhance the current process?
  - (T1) Challenges faced by the organizations
    C1.1: need for further involvement of beneficiaries (and their representatives) in the process;
    C1.2: automation of the recruitment process and its consequences
    C1.3: ICT support to motivate the actors
    C1.4: need for evaluating the success of the innovations
  - (T2) The organization's expectations of the resulting ICT platform and the pilot
    C2.1: automation of the communication between organization and volunteers
    C2.2: management of tools and materials to support the different innovations
    C2.3: involving new actors in the innovation process

The codes were also analyzed in terms of the social innovation process step they relate to, and in terms of "gain and pain" elements (motivation, team building, flexibility of participation, commitment through time, adequate support from coordinators along the process, digital platform support, training material) in the process as highlighted by the interviewees.

Each case study was analyzed individually and then later reviewed by contrasting one with the other based on the core elements of the analysis: the actors; the social innovation process itself and the specific challenges faced by each of these institution's process. This enabled us to trace a common pattern to define the SOCRATIC concept.

4 The Case Studies

The results of each case studies are presented by introducing the organizations, how the data was collected in theirs case studies, their social innovation process (SIP) and the roles played by the different actors in it. Finally, we highlight the aspects of the processes which are of major challenge and could be impacted by the introduction of a new digital social innovation platform.

4.1 CiberVoluntarios Foundation

CiberVoluntarios Foundation (CIB) is a non-for-profit organization whose vision is to exploit information technologies to boost social innovation and to enable citizens' empowerment. CIB was founded in 2005. It coordinates the labour of so called CiberVolunteers, more than 1,500 volunteers all over the world (mostly Spain and Latin America) who actively engage in volunteer activities related to the digital inclusion of populations with low or no access to technology and training. CIB currently has an ICT platform that assists in the management of the volunteers, but it has its limitations. They would like that SOCRATIC to replace their current platform.

The main activities of the foundation are to support, coordinate and leverage a variety of actions that are delivered by the volunteers. These actions include on-site actions, training, courses, webinars and online campaign. CIB activities are usually conducted in cooperation with organisations that support groups of persons excluded from the society, for instance organisations supporting people with reduced physical or mental functionality, people on the verge of poverty, young people with social integration problems, or women facing gender violence.

Participants Involved in the Study

Along with going through the available secondary sources, i.e., CIB website and their training material, the founder of CIB and the volunteers' coordinator were interviewed in order to provide a first understanding of the organization. Based on the first analysis of these interviews and the secondary sources, four user profiles were identified and actors fitting those profiles were interviewed for the case study. Those four profiles consisted of a proactive volunteer, a junior volunteer, a senior volunteer who assumed the responsibility for training other volunteers, and a representative of a beneficiary institution served by CIB.

Roles and Responsibilities

The management at CIB take care of the overarching strategy of the foundation, accounting, fundraising and other organizational aspects. CIB counts with one person dedicated to the management of the volunteers, the volunteers' coordinator. The main tasks of the volunteers' coordinator are to organize the volunteers training, the diffusion of the activities and the setting up of those.

The volunteers play a major role in the SIP supported by CIB, and are considered as the real social innovators. They are usually proactive and dynamic people, who find in CIB an interesting and stimulating way to contribute to society and use their skills for meaningful purposes. Volunteers can choose their grade of involvement depending on their availability and preferences. The flexibility given to the volunteers in relation with their participation in CIB is an important factor for attracting a large number of volunteers.

> *... through my work I had relation with an organization working with disability and this kind of things. So, seeing that there were some options there, through Alejandra, Yolanda, I told them that it would be interesting to work with some collective, even though I wasn't able to go on to give the training because of my schedule... I've seen both... if you're proactive, the organization responds you and you can develop some project... and I have jumped in projects that*

were developing… for me this flexibility/facility is very valuable, if not, you cannot be a volunteer for a lot of time.
Proactive and social innovator volunteer - 04'50

The social challenges that they implement solutions for are determined by the needs of the target social groups, which are represented in the process by beneficiaries organizations. The solutions for those challenges are, in general, proposed by the management at CIB. However, the most proactive volunteers also propose ideas, which are reviewed and approved by the management at CIB. When approved, the ideas are translated into "missions" whose complexity vary. Examples of missions include giving a course in a beneficiary organization or developing a website.

The management at CIB wants to empower volunteers to act "on the field" in order to tackle challenges.

Social Innovation Process

The process starts by the registration and enrolment of the volunteers. New members have to fill their profile. Profiles can then be updated at any time. In the profile, the volunteers share their availability, their preferred activities and information about the population group they would like to work with.

Shortly after registration, the volunteers coordinator contact the new member by phone in order to establish a first contact between the organization and the volunteer. After this call, the new volunteers are invited to receive an initial training aiming at establishing a common understanding of CIB and the SIP at CIB. This step has a legal dimension too: It enables the volunteers to get aware about their rights and duties.

Once the training is fulfilled, the volunteers are allowed to propose and take part in suggested activities. Depending on the kind of activity, the junior volunteer (i.e. less experienced) has to be accompanied by a senior volunteer. This is for instance the case for training activities.

In parallel, the management at CIB proactively identifies social needs which they believe theirs volunteers can help on solving. They enter in contact with beneficiaries organizations with those needs and organize activities to tackle them. The planning of the activities are mainly drawn by the volunteers' coordinator with support of the beneficiaries organizations. As new activities are planned, the volunteer's coordinator gets in touch with the volunteers by mail or by phone in order to inquire about their availability. Simultaneously, the activities are published in their current ICT platform, so that volunteers can manifest their interest to participate. Once the activity is further defined, all the needed information and baseline material for performing the activity is sent to the volunteers. This includes location, previous experiences and relevant knowledge.

In order to carry out the different activities, another important step of the process is the creation of the baseline materials. Those materials are mainly provided by the foundation itself, sometimes with the support of the volunteers. Volunteers support in the elaboration of the materials in two ways. They can enhance existing material with experiences from activities they took part on, and volunteers with special ICT expertise can create new materials. In both cases, the coordinator supervises the creation and management of the materials.

A last step of the process is the evaluation of the activity that has been carried out. The evaluation is performed by the beneficiary institution benefiting from the activity. Although that is seen as an important step, it is currently not done systematically and some activities go without a final evaluation.

Main Challenges

The process at CIB involves a lot of personal contact that is perceived as very important by all parties. The contact between the volunteers' coordinator and the new volunteer adds a necessary human touch to the process and reinforces the engagement of the volunteers.

> *I try to get them motivated… also try to engage them in making us known and so others become a cybervolunteer. I try to motivate them it's very important to motivate them, if not, they fell down (meaning that they drop being a volunteer)… If you are not motivating, if they don't know "there is A".*
>
> *Volunteer's coordinator -01'12"00*

Furthermore, CIB exploits this contact to estimate the motivation level of the newly enrolled volunteer and to establish a common understanding of the expectations, in particular about the importance of the initial training sessions.

Personal contact also applies in the relation between CIB and the beneficiaries organizations. In the cases where proactive volunteers bring the needs of new beneficiaries organizations to the attention of the management at CIB, and suggest a new activity, it is CIB's management who presents the activity to the organizations. This step enforces CIB's institutional responsibility and is understood to establish trust. Other reasons for CIB acting as a link between beneficiaries and the volunteers is the geographic distribution of the volunteers and the lack of implemented ICT support that supports remote collaboration. Thus CIB usually takes care of the prompt and ideation stage and only bring volunteers in the realization of projects.

Personal contact and CIB involvement in the first steps build a large workload over the volunteer's coordinator and the general management of the activities, making it difficult for CIB activities to scale. However, it also cannot be completely replaced by a digital platform. CIB wants, through a new ICT platform for social innovation, to find a way to streamline the process without jeopardizing the trust with members and beneficiaries organizations.

Activities can be proposed by CIB, volunteers or organisations benefiting of solutions to social challenges. The coordinator at CIB encourages volunteers and beneficiaries organizations to proactively suggest activities responding to social challenges.

> *we go through our neighborhood,… and we maybe don't know that there is a little organization with little resources attending people with functional disabilities… and they cannot do it because they're not in the internet, and the one that's not in the internet it seems that they cannot receive help…. a lot know us because of the internet…. In this case the CV will be the mediator or the direct contact with our help. We tell them, observe your neighborhood, your friends,… we tell them to observe around them, because there is always something to do. There is always someone that needs… we give them tips… Is there a little association…*
> *Volunteer's coordinator -21'12*

However, in practice, most activities are proposed by the coordinator. The current ICT platform used by CIB for interacting with volunteers does not stimulate volunteers to be proactive, but rather to respond to the activities proposed by the coordinator. From the management point of view, it is desirable that the organizations representing beneficiaries and the volunteers can take a more actively part of the process. A socially inclusive digital platform can facilitate their involvement in the innovation process.

> *sometimes we get directly in contact with the associations, materialized this help, that this person that answer me the phone, the person that we visit, the person that manage this NGO see this as a resource, a help that will generate wealth for their collective. ...[...] Give them the opportunity for them to grow up, give them this opportunity to get trained... but it doesn't happen do much. ... quite a lot contact us directly, but the one that we are contacting are quite a lot... but when we get into contact to grow our help and permit it gets to everybody, you find this wall, this barrier that say it's not for them. So, it's the only thing that I can find...*
> *Volunteer's coordinator -49'38*

The management of the baseline materials used is different actions is currently a time-consuming activity. A new digital platform should support document management. It would facilitate coordination and cooperation between volunteers.

Finally, the evaluation step can be clearly improved by the introduction of a digital platform supporting the collection of feedback from the involved institutions and for establishing a discussion space where volunteers can share experience and learn from their peers. Those would also help to set key performance indicators (KPIs) related to participation and perceived benefits and allow CIB to better understand their limitations and define measures towards improvements.

4.2 Experts in Team

The "Experts in Team" program (EiT) is a disruptive study program at the Norwegian University of Science and Technology (NTNU) that aims at teaching students teamwork skills and applying their academic competence in an innovative way. EiT is taken by 2,000 students every year, divided in approximately 70 classes (called "villages"). Each village is supervised by a professor, who describes for the village an open ended challenge. In the majority of the cases, this challenge is defined in cooperation with external customers. In each village, the students work in multi-disciplinary teams to solve the customer's challenges. Students are engaged in their projects primarily through experienced-based learning.

Participants Involved in the Study

As in the case of Cibervoluntarios, different members of the EiT program were interviewed in order to collect data to respond to our two main research questions. The same procedure was followed. In a first step, two EiT coordinators were interviewed in order to get a wide understanding of the program. Based on this understanding, two user profiles were identified and actors fitting those profiles were interviewed. The profiles consisted of students currently undertaking the EiT course, and students that had taken the EiT course in previous years. While the first profile allows us to understand the expectations of the SIP and its practice along the course, the second allows us to obtain information of EiT as a lived experience and its impact after the course.

Roles and Responsibilities
The village supervisor (i.e. a professor) and its assistants are trained to guide the students through the whole EiT innovation process. They are responsible for guiding the students in developing team work skills. They conduct different exercises and support the students in each step. Depending on the challenge and the ideas proposed to solve it, the supervisor invites external participants, corresponding to social innovation beneficiaries, to take part in the process.

The students freely propose various ideas and specify their own project implementation plan as long as they stay within the given thematic area and end up delivering the required outcomes. During the whole EiT innovation process, the students, within theirs teams, are responsible to address and strengthen team building and to identify positions for all members in the team. The intention with multi-disciplinarity in the teams is to reproduce the real-work world context.

Social Innovation Process
An important part of the process within EiT is the team building. The aim is to illustrate the benefits of working together from different perspectives, such as better understanding problems and bringing in critical questions in the development of solutions. The EiT program organizes team building sessions where different tools are used. For instance, the "competence triangle" invites the team members to reflect about their competences, both in terms of knowledge and personal capabilities; the "marshmallow and spaghetti tower challenge" stimulates teamwork though finding solutions to building the highest possible tower; the "sociogram" developed by the supervisor reflects the interaction taking place in the team and allows team members to become aware of the dynamic taking place; the "cooperation agreement", designed by the team at the beginning of the course, describes cooperation risks and measures to tackle them, and; the "process report" summarizes and analyses the lessons learnt during the process experiences.

During the development of ideas and solutions, several tools are also used and different issues have to be addressed at different steps following a well-defined calendar. The team has first to describe the problem to be solved, and, within the two first weeks, to propose an idea to solve the problem. To do so, a brainstorming session is organized. Once the team pick an idea and start elaborating it further, a "café dialogue" is organized in order to present the ideas to other teams and exchange feedback. As a third step, once the idea is well defined, the students starts planning and organizing the idea execution. They identify the tasks and the materials needed. Every morning, before the class, teachers are available to give support to the team in case it is needed. Students are graded mainly based on their competence in work as a team, not on the solution delivered.

The target beneficiary group for the solution is more or less involved throughout the process. The course organizes a feedback session with beneficiaries during the ideation period, but apart from that, it is up to the students to involve or not the beneficiaries further in the project.

Main Challenges
The students' assignments are made open as to make it easier for each team to define a project where every member is able to contribute with their particular expertise,

regardless of which study program they come from. It is also is expected to foster student creativity and a strong sense of ownership of the conceived project. However, such open approach can lead to projects that finally don't address the core societal problems of beneficiaries. The participation of beneficiaries is not a major issue for EiT. The course focus is multidisciplinary work and student cooperation in each innovation step. The students are graded on how well they work together, but not on how well their solution address the beneficiaries' problem. Enhancing the communication between students and beneficiaries could increased the value proposition of their project as well as provide a more realistic experience of social innovation. A digital platform has the potential to lower the barrier between students and beneficiaries without reducing the responsibility of the students and their ownership of the project.

> *The supervisors are interested in pursuing steps to make the contact between student teams and beneficiaries as part of the process. But, they don't want to insert many guidelines or requirements which can make the process too narrow. Finding a balance here is one of the key challenges for EiT.*

> *Nonono, they have… this is up to us as a teacher to organize this feedbacks sessions so that we invite… I could even think that we invite CIB to come here and see what they have done… but we cannot say "now we give you a bad mark because you don't…"*
> *Course coordinator 42'51*

Another challenge is about support the projects behind the prototyping phases. The course is too short for allowing enough time to go beyond prototyping. Therefore, students tend not to consider so much the scalability of their ideas or aim at big challenges. Beneficiaries feedback could also help with this aspect. Getting positive feedback and interest from the beneficiaries has the potential of motivating the teams to develop their innovations further even after the course has finished.

4.3 Analysis of Results

The Open Book of Social Innovation [5] SIP definition, presented in Sect. 2, is inline, in terms of steps, with the process undertaken and/or wished by both CIB and EiT. Currently the social innovation activities facilitated by both organizations go through the step of prompts, ideation and prototyping. In the case of EiT, student groups can go beyond prototyping, however that does not happen during the course. Meanwhile, in CiB the current activities are mainly punctual (such as a presentation, giving a course, etc.) and stop at an implementation and replication stage rather than evolving as a separate innovation venture to be sustained, scaled and trigger systemic change. **The steps proposed in the Open Book of Social Innovation are therefore a good starting point for establishing the core SOCRATIC concepts.**

Both organizations would like their supported innovations to go further. CiB would like to facilitate initiatives that go beyond punctual actions and NTNU/EiT would like to provide support for students projects after the course is finished. As a result, **the SOCRATIC concept must go beyond prototyping in order to enable both organizations to help theirs projects reaching bigger social impact.**

Both CIB volunteers and EiT students have little contact with beneficiaries or their representatives. **Enabling participation and collaboration between innovators and beneficiaries across all social innovation stages is a common need for increasing the social impact of the developed solutions.** It would allow offload the workload of CIB, empower both beneficiaries, volunteers and students, and increase the likelihood of designing solutions that match the real needs of the beneficiaries [19].

The introduction of a digital platform has several potential advantages. Community building supported by a platform can contribute to establish trust consolidated nowadays through direct contact with the coordinators at CIB and EiT. If the community achieves an active quorum, the process can become auto-organized and the moderators relieved from the coordination activities.

Flexibility and ownership are key success factors for a digital platform. Volunteers wish to be able to keep control of their engagement and the time assigned to volunteering. Letting teams organize their work and make decisions contribute to motivation. SOCRATIC should therefore define a methodology that supports the innovators, but do not limit their decision power.

In the table below we define generalized requirements towards the SOCRATIC concept and the support level to be obtained via a digital platform. We present those requirements together with the whishes or characteristics of CIB and EiT SIPs.

Generalized requirement	CIB	EiT
RQ1: platform should allow beneficiaries representatives to define the social challenges	CIB would like beneficiaries organizations to pro-actively bring and define social challenges to be solved	Customers present the social challenges to students
RQ2: coordinators provide guidelines and support to participants	CIB contact volunteers and provide an initial training in order to set-up common expectations and understanding	Professors present techniques for fostering teamwork at the student groups
RQ3: coordinators oversees and curate the initiatives they are supporting	CIB wants to approve ideas to be translated to missions CIB wants to oversee the process	Teachers are available to give support in case it is needed
RQ4: facilitating the mapping of skills and interests with ideas and projects	CIB's coordinators currently spend time trying to map and allocate volunteers to actions	Students expertise play a role into which ideas they will be part of developing
RQ5: facilitate discussion and information sharing between participants in the SIP	CIB share baseline information and materials. Volunteers collaborate in the further development of this material	Students present theirs ideas to each other and get feedback through a dialogue café
RQ6: beneficiaries should evaluate and give feedback	CIB would like to improve the evaluation process and	Professors think that initiatives developed

(*continued*)

(*continued*)

Generalized requirement	CIB	EiT
on the initiatives being developed by the innovators	systematically get feedback from beneficiaries	through EiT will have more impact if beneficiaries are more involved
RQ7: the concept model and platform should provide help for the projects to go beyond prototyping	CIB would like to put mechanisms to support initiatives with a longer duration (which stretch into sustaining, scaling and on)	EIT finishes at the prototyping stage but they would like the students to be encouraged and supported to go ahead with theirs innovations
RQ8: the concept model and platform should empower the participants and decrease the organization burden of the coordinators	CIB wants to foster the proactivity of its volunteers and offload the coordinators work	It is intended that students have a lot of freedom to develop theirs ideas but that they receive some guidance as well

5 The SOCRATIC Concept Model

The SOCRATIC SIP follows the same core steps of the Open Book of Social Innovation's SIP [5]. However, the SIP is adapted to:

1. incorporate the organizations' role of facilitating the social innovations;
2. take into account the opportunities (mainly collaborative discussion and production) enabled by a collaborative digital platform, the SOCRATIC platform;
3. be grounded on successful work process coming from the organizations (CIB and EIT) into facilitating social innovation.

In this section, we describe the core concepts of the SOCRATIC SIP and their relationship to the SOCRATIC platform. We start by introducing the different actors of the SIP as they will be referred to in the process steps.

5.1 Roles

The identified roles associated to the Socratic concept model are:

- Challenge Owner: A person who proposes a societal challenge that should be collectively solved. It is a role expected to be played by beneficiaries representatives or a public or private organization interested in solving that societal challenge. In a broader sense, it could be played by any citizen, as long as s-he is willing to support the emerging projects aimed at solve the challenge. It corresponds to EiT's village customers and CIB's beneficiaries representatives.
- Challenge Solver: A person who contributes to solve the societal challenge. The challenge solvers contribute to SIP by helping materialize and realize ideas. Any individual eager to solve a societal challenge and willing to contribute to a social

innovation project can play that role. In CIB, this role is typically played by volunteers and in EiT by the students.

- Challenge Solver-Leader (CSL): A person who takes the lead the development of an idea and/or the realization of a social innovation project. This role is played by the challenge solver driving the innovation and it can be shared and exchanged between the challenge solvers working on the same idea or project.
- Beneficiaries: Members of the societal group targeted by the societal challenge. Beneficiaries may also be actively involved in the process and play a role as a challenge solver or challenge solver leader.
- Coordinators: representatives of the facilitating organizations (such as CIB and EiT). They supervise the SIP and support the different stakeholders involved in the process (Table 1).

Table 1. Mapping roles to actors in CIB and EiT scenarios

Roles	CIB	EiT
Challenge owner	Beneficiaries representatives	Customer
Challenge solver	Volunteer	Student/student's group
Challenge solver-leader	Proactive volunteer	Student group leader
Beneficiaries	Beneficiaries (target groups)	Beneficiaries
Coordinator	Coordinator/cibervoluntarios foundation	Course coordinator/platform moderator

5.2 Process Steps

- Preparation: users need to be "prepared" for SOCRATIC in the sense of understanding the roles of the different actors and the SIP (RQ2). The coordinators take an active role teaching this understanding to users. In order to support them, we defined the concept of Innovation Space. Innovations Spaces are curated spaces (RQ3) in the SOCRATIC platform customized by organizations supporting social innovations, such as EiT and NTNU. There, coordinators can present their vision, introduce the SIP, and theirs role within the SIP in a terminology and context adapted to its contributors.
- Prompts: Challenge Owners describe the social challenge in the SOCRATIC platform (RQ1) and have the opportunity to discuss it with the SOCRATIC community, i.e. beneficiaries, researchers, innovators and professionals of different skills. In this step, challenge owners can get a further understanding of the challenge and create awareness and interest among the community. The platform enables the different stakeholders to collaborate and discuss the challenge. The end of this stage is marked by the final definition of the challenge and the start of the ideation process.
- Ideation: Challenge Solvers brainstorm, collaboratively develop RQ5 and rank ideas that can address the societal challenge through the platform. The whole

SOCRATIC community can participate in this step and innovators should be encouraged to involve the beneficiaries RQ6. This step is completed when the Challenge Owners selects the most promising ideas to advance on the SIP.

- Prototyping: this stage corresponds to the beginning of the ideas materialization. The CSL leading the selected idea, together with other Challenge Solvers, starts a project with the aim of realizing the idea and testing its value proposition. Inspired by the success of agile methodologies in software development [20] and management [21], in SOCRATIC, we advocate that this stage should be performed iteratively through the development and test of minimum viable products (MVPs) with increased complexity. Beneficiaries should test the MVPs and may also participate in co-designing them RQ6. The SOCRATIC platform should facilitate the interaction between Challenge Solvers, Beneficiaries, Challenge owner and other possible investors, and help guiding the Challenge Solver-Leader in adopting an iterative approach following the SOCRATIC methodology. This stage is completed when the MVP provides enough value to the beneficiaries to support its sustainability.

From the prototyping stage, we foresee that the SOCRATIC platform will mainly function as a meeting place between project team and a wider community including the beneficiaries. RQ7 It should allow the project team to present the progress of their project, ask for support and share their learning experiences with the community. The steps after the prototyping correspond to bringing the prototype to market as a solution, scale it and have it socially adopted at large bringing systemic. Those steps are common to the Open Book of Social Innovation's [5] SIP and heavily rely on strategic planning, execution and building networks with customers, partners and policy makers. Much of that traditionally happens outside of ICT platforms, and as of now, we are not yet sure on how they could be supported by the SOCRATIC platform (apart from what is already supported in terms of prototyping.

In this process, the innovation is characterized as an evolving artefact in which the different contributors discuss and collaborate with the support of a digital platform. This artefact starts as a challenge, which once is well defined and published, is open for discussion of potential ideas that may lead to a solution. Once an idea is selected as a viable project, an iterative development project is established with prototypes used to engage with the beneficiaries until the release of the solution.

The role of a digital platform in SOCRATIC is largely to support a wide collaboration among stakeholders along this process and the evolution of this artefact. Such aspect addresses CIB's and EIT's challenge in terms of decentralizing the process and ensuring beneficiaries involvement respectively RQ8.

In practice the collaboration can be facilitated via social computing elements such as web-voting, sharing and conversation threading tailored for each of the social innovation steps. The tailoring of the interface in each step should guide users in following the methodology. So, for example, during the ideation phase, different UI elements will advise users about performing a business analysis, elaborating an elevator pitch and researching the value proposition with real beneficiaries through interviews and other data collection methods. However, to maintain flexibility, the innovators should be able to choose what aspects are relevant and what tools to use. As noticed in both case study, the Challenge solvers will have different time availability and/or skills, making in it necessary to have a flexible approach in order to accommodate them RQ4.

In order to support proactivity and engagement, the digital platform should implement some sort of recommendation system that matches together challenges to users with relevant interests and/or skills RQ8/RQ4. In that way Challenge Solvers could easily identify initiatives they want to collaborate at, offloading the coordinator's work. With the mechanisms in place for Challenge Owners and Solvers to take more ownership of the process, the coordinators would essentially use the platform to supervise the progress and collaborations happening in each innovation.

6 Conclusions and Further Research Steps

The goals of CIB and EiT are different. The two organisations have developed social innovation processes independently from each other, in a pragmatic way, based on experiences gained through voluntarism and training in cooperation work. Despite these differences, the case studies allow us to identify several commonalities between their innovation processes. We also find out that a digital collaboration platform has the potential to facilitate the processes, for instance by releasing the coordinators from a substantial workload, and to enhance the developed solutions by enabling tighter collaboration with beneficiaries. The case studies also a need to educate new volunteers as CIB and students at EiT in implementing social innovation processes. Therefore a digital platform should not solely support collaboration, but also provide methodological guidelines that help a wide audience what social innovation is about and teach them best practices to achieve successful social innovation.

The SOCRATIC project is currently developing a methodology and a digital platform. In our work we tightly work with coordinators, volunteers and students at CIB and EiT. All solutions are developed and tested in an iterative way. Pilots will be conducted to assess both methodology and platform.

Acknowledgements. This work is funded by the EC H2020 CAS project SOCRATIC, G.A. 688228 (http://www.socratic.eu). We thank the participants to the interviews and the project partners for the collaboration during participant recruitment and discussion of the results.

References

1. Simon, J., Millard, J., Lauritzen, J., Carpenter, G., Schimpf, G., Leszek, P.: Doing social innovation. a guide for practitioners. TEPSIE report (2014). http://www.tepsie.eu/images/documents/practitioner_report_final_web.pdf
2. http://www.socialtech.org.uk/nominet-trust-100/2016/
3. Gabriel, M.: Making it big. NESTA report. http://www.nesta.org.uk/publications/making-it-big-strategies-scaling-social-innovations
4. www.socratic.eu
5. Murray, R., Caulier-Grice, J., Mulgan, G.: Open book of social innovation. NESTA report. http://www.nesta.org.uk/publications/open-book-social-innovation
6. www.cap4access.eu

7. www.decarbonet.eu
8. www.litemap.net
9. www.debatehub.net
10. www.kickstarter.com
11. www.fundly.com
12. www.asana.com
13. www.openproject.org
14. www.openideo.com
15. www.quirky.com
16. Hevner, A.R., March, S.T., Park, J.: Design science in information systems research. MIS Q. **28**, 75–105 (2014)
17. Yin, R.: Case Study Research: Design and Methods, 5th edn. Sage Publications, Thousand Oaks (2014)
18. Clare, V., Braun, V.: Teaching thematic analysis: overcoming challenges and developing strategies for effective learning. Psychologist **26**(2), 120–123 (2013)
19. Kujala, S.: User involvement: a review of the benefits and challenges. Behav. Inf. Technol. **22**(1), 1–16 (2003)
20. Schwaber, K., Beedle, M.: Agile Software Development with Scrum, vol. 1. Prentice Hall, Upper Saddle River (2002)
21. Highsmith, J.: Agile Project Management: Creating Innovative Products. Pearson Education, London (2009)

Connecting Citizens: Designing for Data Collection and Dissemination in the Smart City

Donald McMillan(✉)

Stockholm University, Stockholm, Sweden
donald.mcmillan@dsv.su.se

Abstract. This paper presents two case studies of citizen data collection and dissemination applications, developed for or by three different local authorities in Northern Europe. These case studies highlight the challenges in meeting the goals of Open Data, of involving citizens as sources of information, and of engendering and maintaining trust as a service provider all at the same time. The challenge of making data open can be seen as at odds with protecting the privacy and safety of citizens when it is sourced directly or indirectly from their actions. Encouraging citizens to collect, curate, and submit data can create misguided expectations of influence over the processes of local government, and disillusionment where action or feedback are not forthcoming. A local authority is trusted to provide information that is verified and for which it is accountable. Balancing this with goal of disseminating the results of citizen sourced data collection activities can result in frustration for developers, users, and local authority employees. In response to these issues this paper presents the following four design opportunities: *probabilistic and personalised representations of data, making accountable the use of collected data, respecting the boundaries of data,* and designing for the *graceful degradation of resources*.

Keywords: Smart city · Big data · Open data · Crowdsourced data

1 Introduction

The conflicting visions of what a modern city should provide for its citizens, and how that provision should be decided upon, targeted, and measured can be seen not only between cities and states, but between and within departments in modern cities. Understanding the differing interpretations of what it means to be ‘smart’ and what the path to securing a place on the Forbes list of Top Ten Smart Cities is key to being able to understand the behaviour of local municipalities with respect to innovative and disruptive technology.

There has been considerable interest in the opportunities for better data management and collection to improve city functioning. More ambitiously this has been at the heart of arguments for the “smart cities” program, attempts to revolutionise the functioning of city life through technologies of different sorts. Yet these discussions seem often strangely disconnected from the actual practicalities of city activity. Cities governments are still often bureaucratic organizations, with complex leadership structures and conservative attitudes to innovation. Moreover, IT development and management

I. Kompatsiaris et al. (Eds.): INSCI 2017, LNCS 10673, pp. 119–131, 2017.
https://doi.org/10.1007/978-3-319-70284-1_10

is often contracted out to private companies, making for complex relationships and incentives for those involved. Indeed, some of the ‘smart city’ rhetoric is perhaps more about corporations attempting to lock cities into particular technological platforms and vendors.

This paper builds upon previous work in HCI understanding the expectations, and desires, and constraints on the stakeholders involved on all sides of city procurement of new technology [15]. Following on from this work, this paper is structured around two case studies. The first is of a community generated map of addiction and food services, which we use to highlight the challenges of trust, provenance, and branding in the flow of data. The second covers a set of similar, yet independent, applications developed in different cities to support and quantify the increase in cycling as both a health and environmental benefit to citizens. Using these case studies, this paper presents opportunities for design with and for the city fitting with the real-world constraints exposed here and in previous work.

2 Background

Smart City programs typically seek to understand, manage, and improve city functions, often with a top-down approach, through the use of distributed sensing technology and data processing of various kinds. The production of sophisticated data analytics for understanding, monitoring, regulating and planning the city is a key issue underlying the idea of “smart cities” [11] – a vision “for stimulating and supporting innovation and economic growth, and providing sustainable and efficient urban management and development.”

Townsend [17] emphasises the role of ubiquitous digital technology in improving how cities function and operate. He describes smart cities “as places where information technology is combined with infrastructure, architecture, everyday objects, and even our bodies to address social, economic, and environmental problems.” Yet, he also questions the motivations underlying attempts to transform cities, arguing that “[l]ooking smart, perhaps even more than actually being smart, is crucial to competing in today’s global economy”, and that such optics may, in fact, be “the real force driving mayors into the arms of engineers.” Goldsmith and Crawford [7] point to another challenge cities face in the race to adapt their operations to ‘smart’ ways of working with technology, where legislation and a rule-bound approach to government becomes an obstacle when working with the vast amounts of data generated by the same technology and by the citizens they serve. They argue for a focus on results rather than compliance to legislation, and on problem solving that combines the city’s data with collective knowledge and data generated by the citizens.

Kitchin [11] sees exciting opportunities in cities’ move to make use of new data streams to help both governments and citizens to make sense of the city, but he points to serious concerns about “the real-time city”, too. These include the politics of big urban data, technocratic governance and city development, the corporatisation of city management and the risk of technological lock-in, system vulnerabilities, and, finally, ethical issues involved in surveillance and control. Building on the “smart city” critiques by Greenfield [8] and Townsend [17] that advocate for grounded approaches that account for civic everyday realities, Kitchin [11] points out the risk that, without critical interrogation, future "smart cities" may fail to reflect the desires of wider society and, instead, prioritize narrow corporate and state visions.

HCI researchers have presented a range of services that work in the interface between citizens and the city, such as tools for civic engagement that enable citizens to engage with city governance. These include, for example, tools that support civic involvement in democratic processes through situated voting, or invite dialogue around politics of place making [5, 12, 16], and technologies that inform city traffic infrastructure and environmental services with the help of mobile sensing and crowdsourced data [1, 9, 13, 18]. Others have proposed systems for improving the accessibility of the city's services and data, for instance, by supporting interactions between families and case workers when planning parental leave [3, 4], by presenting local crime data on mobile devices to reduce fear of crime [2] or by encouraging smarter water consumption [6]. While this body of work does identify complexities in working with data in city administrations, the main emphasis has been on developing and demonstrating new technologies.

Of particular relevance to our present study are two further pieces that address broader challenges in deploying technology within the public sector: First, Le Dantec and Edwards [14] performed a year-long ethnographic study of ICT use in the public sector, looking specifically at ICT projects that cross institutional boundaries. The authors argue that crossing such boundaries is, in fact, a central part of the work in the public sector. They use the notion of *scale* to describe these boundaries and the complexities they present at different levels: *"(…) cooperative systems with large numbers of users (across independent organizations), and long lifespans (as tools for enacting public policy), and whose use encompasses communities that cross local, regional, and national contexts"*. Second, Harding *et al.* [10] observe that the perceived value and sustained use of technologies for civic engagement has remained low. They argue that prior work has been, perhaps surprisingly, too citizen-centered, and has, as such, failed to account for the needs and concerns of civic authorities whose responsibility it is to ensure the accountability of the produced data. One conclusion from this work is that interaction between the different stakeholders is needed to overcome mistrust and lack of appreciation of the challenges other actors face. They point, specifically, to three key aspects to be considered in design processes: (1) how authorities' IT systems are opened to enable new activities with new forms of data, (2) political and organisational factors hindering transparency, and (3) changes that are needed to work practices within cities to support the development and use of new tools.

3 Method

This paper builds upon a close working relationship with a local city authority in a north European capital city, involving collaborative research projects where the city has provided forums for research, contact to participants, and real word problems to be tackled in collaboration with different city authorities, particularly together with the city's youth services department. Following on from this, we expanded our perspective and collected research materials with a wider scope. We conducted 20 interviews across five northern Europe cities – interviewing city management, IT managers, contractors and developers. Seven interviewees were managers or staff working in the city government either in front line positions or managing projects that made use of data and software. The other thirteen represented third party developer companies, civic initiatives, and non-profits.

The interviewees were selected on the basis that they had key roles in the development of publicly released citizen-facing applications developed directly by the city, in collaboration with the city, and/or commissioned by the city. The targeted applications had to meet the selection criteria that they were (1) data-centric and citizen-facing, either providing citizens access to city data or producing data to be used in city government in some way, or both, (2) past the project phases of funding, planning and development, and to the point of launch (although not all projects launched). By approaching people in different roles in such projects, the aim was to capture a range of perspectives of how data-driven applications are born, developed and maintained, from planning and strategic decisions to building and practical maintenance.

The interviews were semi-structured and lasted between 40 min and 1 h 20 min. The majority of the interviews were conducted in English. In eight instances, interviews took place in the interviewee's native language (other than English). In these cases, the quotes we present have been translated verbatim to English. The interviews were informed in part by our experiences in working with our city over a period of seven years. This served as an informative background for our data analysis rather than providing specific data for this study per se. The goal was not to find statistically generalizable points, but rather to generate concepts and understandings for working with local authorities. As such, the approach to analysing the interviews drew on an interpretivist stance, with the development of an understanding of the problems and practices of those being studied. The analysis involved open coding of the interviews, and the development of themes through an iterative process of concept development.

4 Future Maps

This section describes the lifecycle of the Future Maps project: a community based project to crowd source local expertise on the availability of services for addiction and food services such as recovery meetings and foodbanks. After documenting the birth and lifecycle of this application this section concludes with design recommendations for other crowdsourcing projects with a goal of integrating with local authorities.

4.1 The Future Maps System

Leveraging community created, local knowledge is one major goal of Smart City projects and the Future Maps system developed as part of the Glasgow Future Cities Initiative was built around this. The project started by recognising that there was a wealth of information in the community around services that support addicts on their path to recovery collected by a charity in a particularly deprived area of East Glasgow.

"They have a map of the services that their users can access. And they've done that map. The problem was it was on somebody's laptop so if he wasn't in the office then nobody knew where to go."

There were a larger number of services to help those recovering from addiction, either official from bodies other than the city authority or services that could be used as such like sports groups or craft classes, that were not displayed for that purpose on the official council, health service, or policing websites.

"There might only be two or three things on the council website or mentioned by the doctor but in reality there is much, much more."

Another issue with the current official data was that it didn't conform to the socio-cultural realities of living in the areas where the services were being provided. The social structure of some parts of the city means that the borders between neighbouring areas are unsafe to cross by residents, and this segregates the services provided by the city based on the geographic territory that they fall into.

"For example DC [sic] and MH are two different territories and people don't go from DC to MH. And MH to DC...So to the city, they've put a swimming pool where every inhabitant has a park within X meters of them. Actually, that's not necessarily true because it might be outside of their boundary."

By providing the local citizens a way to map what they saw as the services available to them, regardless of physical geography, this was a way to provide meaningful data for those in that area and to highlight shortcomings in provision.

Each of the data points initially was of an activity that would be available to recovering addicts as a way to help them stay on track on their recovery.

"If you're going to a recovery café for example, and you can't find it because it's in a big complex then you're just going to the pub and you've made more damage then."

Those providing the data were therefore careful to give detailed descriptions of when, where, and how to get to the activities that had been mapped. They also included photos to help guide the users to the correct place.

But the focus solely on addiction services was something that was seen as a problem of sorts, by those adding to the system and by the city authority. To tackle the social and political issues around this the categories of services on offer were broadened, this also introduced users of services in one category or subcategory to the services available in the others.

"Drug addiction, alcohol addiction and then we had food banks and that made a pretty bleak picture. These are the services that people need. So we made addiction services and made subcategories and then food we had food, we made it food and growing so the food bank and the growing spaces so that they could be the two ... And for somebody as well at first they're not going straight for the drug category. They can do something so it changes the perception."

At this point the solution to the problem was technically complete, but there was a lack of technical skill in the community to keep the data updated. This was addressed with a series of training workshops, a manual, and a video. This also had a knock on effect of raising the general technical literacy of some of those involved.

"One of them kept forgetting his password. Because he didn't have an e-mail address he couldn't reset it so somebody showed him how to use an e-mail address and saying at the same time well you'll need one anyway to look for a job."

After training the solution was usable and, to some extent, sustainable for the area in which it was deployed. It was also becoming more widely known within the city authority and within the community of interrelated charity organizations which provide support services for different groups across the city.

4.2 Issues with Future Maps

There were, however, a number of issues with the Future Maps service. Maintenance and expansion were up against a number of barriers. Expanding the scope of the map itself, by integrating other data such as events at sports centres or events at health centres, ran into the problem of boundaries. As discussed in detail in [15], data – and the access to data – is a highly charged issue within local authorities, with the desire for openness directly competing with the safety and control that hoarding data provides. These boundaries are also felt between state organisations, so while the city authority may compile and own data regarding the sports centres in the city, the data on what health and wellness classes are provided at the facilities run by the national health service are separate. These organisations were not willing to readily provide the information to be presented in a way outwith their control, with the express concern that this would have the possibility to tarnish the trust the citizens had in them. Data from other sources, especially crowdsourced data from citizens, presented alongside their official data could be, in their eyes, misconstrued as being official or at least verified by the city or the health authority – meaning any mistakes in that data would reflect badly on them.

Another issue with expansion was funding. This was wholly supported by a charity that received donations expressly to work with addiction in one area of the city – they had neither the resources, nor the ethical leeway in the use of their funds, to support other areas of the city and refugee services when asked by another charity. The city authority was also unable to directly fund expansion of the service as it would then become part of the service provision of the city and fall under government procurement regulations.

Without going through procurement the small agency, which through a small innovation fund was able to work with digitising the map, was unable to provide the manpower necessary to take it further. The charity was able to fund server costs, and saw this as a useful part of the service they provided. However, the changing economic climate and the continual churn of volunteers means that updating and verifying the information became a burden forcing them to, hopefully temporarily, make it only available at their main centre where someone would be available to monitor what data was being accessed and make sure that it matched what they knew of the ever shifting landscape of provision available to those they help.

4.3 Design Lessons from Future Maps

While [15] focused on learning to understand the complex organisational structure within which such applications would find themselves, this paper instead attempts to work within the boundaries and with the conflicting goals of the actors present in local authority development.

The first design opportunity here would be to respect the organisational boundaries in the data. While there is a value to aggregating data from a number of different sources to attempt to present a complete picture this is not the only way to present a number of data sources in order to provide users with all the information that they need. There are a number of ways to provide clear separation in the UI. One is using map layers and different icon sets and providing the data providers the opportunity approve which other data sources their data would be able to be shown alongside (meaning that on selecting a course that was not approved, the data layer would be removed from view). In the case of Future Maps, designing to respect such boundaries may have provided the technical, as well as the political, separation necessary to allow the charity focused on addiction to open the service for other charities with other responsibilities.

The second design opportunity presented by the Future Maps service is to allow for graceful degradation of resources. While graceful degradation is a common design principle with respect to device capabilities, meaning that a website, for example, would still provide a good user experience on a device that did not support functionality the main design takes advantage of. In the case of Future Maps, designing to support the waxing and waning of resources could take a number of forms. One would be to invest significant effort in providing stand alone training resources that would allow the service to more easily changes in administration users without an overlap between them. Another would be to ensure that any online-only services would come with a pre-built solution to allow them to be moved seamlessly between paid-for hosting and to be run on a single computer or laptop in-house. On the user facing side, it would also be important to make visible the age of any data presented and provide the ability for users or admin staff to verify data points – with the same UI features providing accountability and age for the verification.

5 Two Cycling Apps

We are able to compare and contrast two projects with very similar goals across two cities in northern Europe. Both services share the fundamental idea that citizen-sourced data from cycling can be leveraged to inform and improve decision-making on city planning around transport – for the betterment of cyclist and to encourage a healthier lifestyle in each respective city. They are not directed at cycling enthusiasts – instead of providing feedback on performance they invite cyclists to contribute their route data to be analysed. The two apps, although they differ on the surface, have very similar functionalities; they use GPS to enable the user to track their ride from start to end point, calculate distance, average speed and total riding time. This data is then submitted for analysis as a service to the city:

"What we provide is digested information out of this, (...) a report or an analysis with some recommendations, perhaps how to improve cycle planning expenditure."

Although both were successful in gaining users, neither of them have yet succeeded in the goal of influencing the city in a meaningful way. This section draws on their respective stories to illustrate the challenges in working with the city.

BR is an example of an initiative where an external actor has initiated and developed a service in collaboration with the city. As the founder explains:

"I'd been thinking about this problem with different modes of traffic for quite a long time, ... it is not really working very well, the cycle planning in part within the transport planning ... I was like, "Something needs to be done.""

The traffic authorities agreed with the problem formulation and entered into a collaboration to make the service a reality, but over a year after launch the city is yet to use the resulting analysis the way it was intended and effect real change.

"We think that you can spend ... tax payer money much more wisely is actually saying to the officers, transport officers, "Actually, you have to change a bit about how you do this and you think about this. ... So that's basically a challenge, to make them pay to do that, to change" – S2.

This illustrates how just providing data is often not enough to affect positive change within the city. Data ties in with city functions and the processes that make up these functions have to be changed to accommodate any change in the data.

GC, on the other hand, was born from the same Future Cities funding initiative as the Future Maps example above. Although seemingly more straight forward in terms of funding and approval from within the city, the full concept of the service was not actually better grounded. The idea of feeding cycling data back into the city came from the developers themselves:

"Yeah, so the spec we have is kind of one paragraph, a cycling app, and that's it. Map usage, cycle usage and we've been at it and we've done what we thought is best."

This meant that the project had similar challenges as the BR app when it came to making the relevant administrations in the city understand the value of the data and achieving the type of impact on decision-making intended in the design of the service. The flow of data produced by the GC app presented a different set of problems:

"(...) we had to work within privacy impact and stuff within the council, be compliant with the Information Commissioner's Office. They would check and double check and treble check and that really slowed progress."

The city was eager to aid in the production of an app that encouraged cycling, but when they realized that data would be collected from citizens and then shared without being manually checked by a city employee, the eventual launch was delayed by over nine months. The developer interviewed for this study felt the city's concerns were unwarranted.

"I think there was a culture of fear around what's going to happen. What if? What if? What if? Which is, again, its insane to have projects that you have constantly fighting against what ifs rather than doing it and seeing."

To date, neither city has used the data produced by the cycling services in any significant way. The developer of GC expressed frustration that what was discussed when the app was being developed did not match the actions he saw once it was eventually released.

> *"In all the press releases that was the line they went with. I would like to see it happen. I wouldn't hold my breath on it. If the council doesn't use it it's not like other people will."*

5.1 Issues with the Cycling Apps

The two main issues that will be focused on here are the collection and dissemination of potentially identifiable citizen data, and the challenge of influencing city policy with citizen data.

The first issue, that of the risk averse nature of local authorities when dealing with the confluence of new technology and potentially identifiable citizen sourced data brought under their control, is understandable. As noted above, local authorities place a lot of value in the trust their citizens have in them and are protective of that trust. Anything that has the possibility of eroding that trust has to have clear and overriding benefits in order to be successful within the local authority. The fear of data, either leaked or officially released, causing real harm is an ever present worry for those tasked with managing open data provision in city government [15]. In this case, the developer of GC recounted a fear put forward to them of a route being seen to start near a hospital and end near a residential area being reasoned by a person with malicious intent to be the commuting route of a nurse or doctor. This fictional actor would then be able to lay in wait for the health worker to assault them on their way home. In a similar vein, the head of the agency working with Future Maps recounted a problem they had with using user stories to explain how the system worked to those outside the charity organisation. They had used a common first name and a fictional street name to describe each of the fictional users to be used as examples in showing how to use the service, but this was vetoed by the local authority because of the potential reputational harm it may have to someone of the same first name in a similarly named street – if they existed – to link them to addiction services.

Both of these interactions with the local authority help paint the picture of an understanding of data dissemination, its power, and the limited control that the authority has over it which – while not unreasonable – proves to be a barrier for development.

The second issue highlighted by the cycling apps is the challenge of influencing city policy. As discussed in previous work [15], this is partially an issue of differing temporal scales between application developers and city planning and partially to do with a lack of clear understanding of where and how the city would incorporate such citizen sourced data into its planning process. So the issue here is two fold, on one hand there is the issue of providing the data in a format, time, and place for it to be influential to the city and on the other hand there is the issue of making visible the influence that the data has to those providing or curating it.

5.2 Design Lessons from the Cycling Apps

The issues highlighted above could be dealt with by working with all stakeholders to ensure a shared understanding of the goals, risks, invested value, and processes involved [15]. However, each of these issues also presents opportunities for design innovations to meet the challenges from which they are born.

The issues around sharing data that has the possibility to be identifiable or attributable to any citizens can be mitigated by designing probabilistic and personalised representations of the data for each user. In the case of the bike routes, the solution to obfuscating the data was to trim the start and end of each cycle ride in order to make it harder to point to an individual residence or workplace. The issues raised by those in the local authority were, in some ways, valid in that such a solution does not protect users on less well used paths nor those that work or live in sparse areas of the map. A more robust solution would be to provide a probabilistic reimagining of routes as they are shared publicly. In such a display the probability of a cycle route following a particular road would take into account possible other routes on the map, the number of unique users who cycle that route, and the frequency with which that route intersects other detected routes. This would mean that a cycle route off the beaten track used by one person to commute would be unlikely to be shown to others, a less well used route that crossed one or more common routes would be shown with deliberate errors, and the most common routes would be shown without modification. The personalisation of the display would be used to show routes that the user has cycled along in their full fidelity irrespective of the number of others that have used that route.

Designing with the goal of influencing city function is more challenging. One aspect of the frustration felt by the developers, which was posited as a reason that the use of these city-focused cycle tracking applications may drop off faster than the commercial, performance focused competitors, was the feeling that the data was left unused by the city employees that it was presented to. In order to make the use of this data accountable, the developer could focus resources on the interface provided to the city employees and develop methods to explore and interrogate the data. Providing functionality in this service for the city official to mark routes as important, or to export the results of a certain enquiry in a format easily incorporated into internal reports and presentations, would encourage the city employees to work within the system. This use could also be used to provide feedback to the end users from whom the data was sourced, showing them which routes had been receiving the most attention from the city authority and engendering a sense of purpose and progression in the use of the app even if new city infrastructure was still at the pre-planning phase. This mechanism could also be used for the city to announce when projects had been approved directly to those who would feel the impact the most rather than through the standard practice of public yet obscure planning meetings.

6 Discussion and Conclusion

This paper has used two case studies drawn from a larger body of interview based work focusing on technological innovation in cites to provide concrete design opportunities available to developers to mitigate some of the most pressing issues when working in such an environment.

In recommending that developers design to *Respect Boundaries of Data* this paper presents a pragmatic solution to the problems that these boundaries can cause. Breaking down such boundaries may be a worthy goal, but beyond what can be reasonably expected from any development team working within their confines. Designs that respect and expose such boundaries may over time, by raising awareness of the boundaries to citizens and municipal employees not directly faced by them on a day to day basis, erode them to an extent where such a recommendation is not necessary.

Designing for the *graceful degradation of resources* is an important counterpoint to the idea of a governmental contract being a golden ticket. Due to the complex nature of procurement at the local authority level most of the applications we discussed with our interviewees, and all of the ones described in this paper, were not developed to be part of the main service provision of the local authority. They were mostly funded as one-off expenses through innovation budgets of one sort or another. While there was the expectation on the part of many of the developers that a successful system would have the chance to become part of the service provision [15], the only examples of this were the cycling apps where one was brought in-house to be managed by the developers already employed by the authority and the other was spun out as a company to sell data analytics back to the authority as a service. Understanding the chance of limited or changing resources assigned to their application allows developers to design for genuine longevity and continued engagement, giving those services a greater chance of being picked up as part of service provision in the future. In this vein, it is possible to start such projects with the goal of becoming part of the service provision. However, in doing so any design must be done with a clear understanding of the general procurement procedure and any funding specific constraints or opportunities resultant from the innovation funding. Providing a clear path to release and a clear benefit for the citizens may not be enough to secure the projects continuation, or to stop another company bidding to provide a competing system if the value is demonstrated.

Designing *probabilistic and personalised representations of data* is one solution to managing the sensitivity that city authorities have to the collection and dissemination of potentially identifiable information from citizens, but it also provides a general tool for the development of more privacy sensitive crowdsourcing applications.

Designing to *make accountable the use of collected data* can be seen as an opportunity to support and encourage the integration of crowdsourcing applications into city functions while allowing the city to present itself as more responsive and connected to the citizens under its purview.

Taken together these four design principles, based on real world examples of issues faced by developers working with local authorities, can be viewed as opportunities to better develop applications that fit with the realities of city organisations. They can benefit not only commercial developers, but also researchers using the city as a testbed.

As cities increasingly move towards introducing technology and data driven solutions as part of their service provision, more and more small scale research and private exploratory projects can be expected to be funded to explore the options available to them. This paper provides concrete recommendations for these projects going forward.

References

1. Aoki, P.M., Honicky, R.J., Mainwaring, A., Myers, C., Paulos, E., Subramanian, S., Woodruff, A.: A vehicle for research: using street sweepers to explore the landscape of environmental community action (2009). https://doi.org/10.1145/1518701.1518762
2. Blom, J., Viswanathan, D., Spasojevic, M., Go, J., Acharya, K., Ahonius, R.: Fear and the city: role of mobile services in harnessing safety and security in urban use contexts. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 2010), pp. 1841–1850 (2010). https://doi.org/10.1145/1753326.1753602
3. Bohøj, M., Borchorst, N.G., Bouvin, N.O., Bødker, S., Zander, P.-O.: Timeline collaboration. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 2010), pp. 523–532 (2010). https://doi.org/10.1145/1753326.1753404
4. Borchorst, N.G., Bødker, S., Zander, P.-O.: The boundaries of participatory citizenship. In: Wagner, I., Tellioğlu, H., Balka, E., Simone, C., Ciolfi, L. (eds.) ECSCW 2009, pp. 1–20. Springer, London (2009). doi:10.1007/978-1-84882-854-4_1
5. Crivellaro, C., Comber, R., Dade-Robertson, M., Bowen, S.J., Wright, P.C., Olivier, P.: Contesting the city: enacting the political through digitally supported urban walks. In: Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI 2015), pp. 2853–2862 (2015). https://doi.org/10.1145/2702123.2702176
6. Erickson, T., Podlaseck, M., Sahu, S., Dai, J.D., Chao, T., Naphade, M.: The dubuque water portal: evaluation of the uptake, use and impact of residential water consumption feedback. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 2012), pp. 675–684 (2012). https://doi.org/10.1145/2207676.2207772
7. Goldsmith, S., Crawford, S.: The Responsive City: Engaging Communities Through Data-Smart Governance. Wiley, Hoboken (2014)
8. Greenfield, A.: Against the smart city. Do projects (2013)
9. Hara, K., Le, V., Froehlich, J.: Combining crowdsourcing and google street view to identify street-level accessibility problems. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 2013), pp. 631–640 (2013). https://doi.org/10.1145/2470654.2470744
10. Harding, M., Knowles, B., Davies, N., Rouncefield, M.: HCI, civic engagement and trust. In: Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI 2015), pp. 2833–2842 (2015). https://doi.org/10.1145/2702123.2702255
11. Kitchin, R.: The real-time city? Big data and smart urbanism. GeoJournal **79**(1), 1–14 (2014). doi:10.1007/s10708-013-9516-8
12. Korn, M., Back, J.: Talking it further: from feelings and memories to civic discussions in and about places. In: Proceedings of the 7th Nordic Conference on Human-Computer Interaction: Making Sense Through Design (NordiCHI 2012), pp. 189–198 (2012). https://doi.org/10.1145/2399016.2399046
13. Le Dantec, C.A., Asad, M., Misra, A., Watkins, K.E.: Planning with crowdsourced data: rhetoric and representation in transportation planning. In: Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW 2015), pp. 1717–1727 (2015). https://doi.org/10.1145/2675133.2675212

14. Le Dantec, C.A., Edwards, W.K.: Across boundaries of influence and accountability: the multiple scales of public sector information systems. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 2010), pp. 113–122 (2010). https://doi.org/10.1145/1753326.1753345
15. McMillan, D., Engström, A., Lampinen, A., Brown, B.: Data and the city. In: Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI 2016), pp. 2933–2944 (2016). https://doi.org/10.1145/2858036.2858434
16. Taylor, A.S., Lindley, S., Regan, T., Sweeney, D., Vlachokyriakos, V., Grainger, L., Lingel, J.: Data-in-place: thinking through the relations between data and community. In: Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI 2015), pp. 2863–2872 (2015). https://doi.org/10.1145/2702123.2702558
17. Townsend, A.M.: Smart Cities: Big Data, Civic Hackers, and the Quest for a New Utopia. W. W. Norton & Company, New York City (2013)
18. Yoo, D., Zimmerman, J., Steinfeld, A., Tomasic, A.: Understanding the space for co-design in riders' interactions with a transit service. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 2010), pp. 1797–1806 (2010). https://doi.org/10.1145/1753326.1753596

Politicians Driving Online Discussions: Are Institutionalized Influencers Top Twitter Users?

Anna S. Smoliarova(✉), Svetlana S. Bodrunova, and Ivan S. Blekanov

St. Petersburg State University, Universitetskaya nab. 7-9, 199034 Saint Petersburg, Russia
a.smolyarova@spbu.ru

Abstract. Embeddedness of politicians and political organizations in a discussion defines its level of institutionalization and creates a public arena for collaboration between publics and institutional actors. Thus, testing whether traditional hierarchies (in terms of presence of politically institutionalized actors) show up in online discussions deserves scholarly research. Moreover, it is also important to see whether more democratic societies show patterns of public involvement of politically institutionalized users that would differ from those in more authoritarian contexts.

To assess the 'influencer' status of politically institutionalized actors on Twitter cross-culturally, we have selected conflictual Twitter discussions in Germany, the USA, and Russia, all based on violent inter-ethnic clashes. Using vocabulary-based web crawling, we collected data on them and formed samples of top users selected by four activity metrics and five network metrics, to assess the positions of political users in the top lists and correlations of user status with their top list ranks. To this, we added qualitative assessment of presence of political users in comparative perspective.

Our results show that, in all the cases, presence of political actors in online discussions is scarce; also, political actors tend to fail to link user groups or stay in the center of discussion. There is also meaningful divergence of Russia from the pattern that Germany and the USA show: while in these countries politicians gain user attention based on content, in Russia it is the status itself that matters, and political users tend to gain weight in the discussion structure despite low attention levels.

Keywords: Twitter · Influencers · Political actors · Web crawling · User metrics

1 Introduction

Twitter has been often described in the literature as a user-driven platform where horizontality of user connections puts into question traditional hierarchies of actors who influence the debate; e.g., Twitter dramatically individualized participation in collective actions and social movements [1, 2]. According to the studies of digital

I. Kompatsiaris et al. (Eds.): INSCI 2017, LNCS 10673, pp. 132–147, 2017.
https://doi.org/10.1007/978-3-319-70284-1_11

protest and ‘hashtag activism’, Twitter has enabled longitudinal campaigning that changes structure and content of public discussions [3]. Case studies of Twitter discussions on selected conflicts paid more attention to the content of debate and the character of public communication rather than users participating in the discussion [3, 4]. Also, several important works focused on Twitter activities and strategies of single categories of users – predominantly communicative elites, like politicians and journalists [cf. 5, 6]. But there is still lack of research on what place such users occupy in *ad hoc* discussions on urgent public matters, whether they regain user attention and authority in the new communicative environment, whether they occupy positions that help organize meaningful discussions, and what lies behind their popularity.

Moreover, there is no consensus among researchers on what kind of users becomes influential and whom to consider an influencer at all, as well as which user metrics one needs to look at to collect a sample of influencers. Our earlier research deals with this issue and offers a framework for selecting top users [7, 8]. In this paper, we apply this approach to assess the political institutionalization of online discussions in cross-national perspective – that task is, again, under-addressed in current research on social media platforms.

Today’s proliferation of social media comes across cultures and political regimes, and it poses another research question – whether national traditions of public discussions (or at least some of their features) show up on similar occasions online, or the platform features and limitations make them irrelevant. In this paper, we address this question by comparatively assessing the discussions in Germany (a European democracy), the USA (a non-European democracy), and Russia which shows authoritarian trends in media development, along with other areas [9–13].

In search of comparable discussions, we looked at three cases of inter-ethnic conflicts that gained national attention and triggered extensive networks of Twitter interactions: the Cologne mass harassment of 2016, Germany, the Ferguson unrest of 2014, USA, and the Biryulyovo bashings of 2013, Russia. All these cases were among national Twitter trending topics and all of them caused nation-wide discussion and social polarization.

For all the cases, we assessed the extent to which the offline user status may change the potential of becoming a Twitter influencer (top user by several parameters); we also looked at the nature of the users that were present in the discussions. In Sect. 2, we present the research framework and literature overview; Sect. 3 provides our research hypotheses; Sect. 4 is dedicated to the methodology and data collection; Sect. 5 presents and discusses our research findings.

2 Influencers in a Horizontal Network: User Status vs. User Rank in Twitter Discussions

2.1 Twitter: A Platform for Reciprocal Communication?

The network theory applied in the Twitter studies draws attention to the structure of the Twitter discussion based on weak links and cluster dynamics, the formation of horizontal influencers and various approaches to their analysis, as well as the relationship of the structure and content of tweets in their relation to the type of user.

In comparison to other social networks, Twitter is often characterized by nonlinear dissemination of information [14] and low level of reciprocal relationships [15–17] creating of Twitter a picture of quite a loose network of non-reciprocated communication. At the same time, a lot of research has focused on finding communicative disparities and hierarchies of various kinds in this seemingly equal and loose discursive milieu.

Despite the initial horizontal character of Twitter communication, scholars revealed several patterns of hierarchy in relation to user influencer status.

First, users differ in terms of activity: 'the more active Twitter users frequently engage in gathering and sharing what they perceive to be relevant materials, for example tweeting links to further information (or retweeting relevant posts of other users) to their own followers or to Twitter communities formed around topical hashtags' [14]. Other authors [18] found that users with higher level of message communication have stronger and more selective influence on the information in Twitter.

Second, although Twitter is seen by scholars as 'a platform that works on the idea of reciprocation' [4], they have found various types of user inequality in the patterns of reciprocation. Thus, it has been shown that users of one type (for e.g. Twitter accounts of news organizations) usually gain attention from users of the same type (other news organizations) [19] rather than expand their network to ordinary users. Follower-followee relationships on Twitter may form a kind of relatively stable, long-term network that tends to represent the long-term interests of users [14].

Third, offline group representation patterns may influence differences in user impact upon how the discussion goes. For instance, as minorities are less connected to social movements, they get lower visibility among the influencers [20–22]; or, 'young people from historically underrepresented communities neither experience themselves as members of a singular community on- or off-line, nor uniformly choose to follow or interact online with elites who have emerged as political leaders' [21]. Nevertheless, the optimistic view that is shared by a large group of scholars says that Twitter enables communities by 'artefacts of engagement - comments, photos, videos, tweets, and news stories' [21: p. 239] which, as the same author claims after studying involvement of young people into Ferguson events and online discussions, "provided an avenue into what Dahlgren called the 'proto-political'" [23, 21: p. 244]. Participation in online discussions made this social group also visible to other parts of the community; this was 'important, as both strong collective identities and strongly felt shared grievances are known to play a significant role in spurring members of one's extended social networks to consider participating for the first time' [21: p. 244]. Clark also identifies 'Black Twitter' as 'a space in which counter-publics may form as people find and follow one another, engage in discussions about the meanings of Blackness, and discuss strategies for engaging in political action that arise from those meanings' [21: p. 239]. But at the same time, this implies that 'Black Twitter' in some respect opposes the wider discussion environment and lacks political institutionalization.

Fourth, a lot of research is focused on how the platform features of Twitter and structural peculiarities of the discussions bring certain users to the positions of influencers. The notions of gateway formation [18] and gatewatching [24, 25] demonstrate that networking patterns strongly influence the potential discussion hierarchies. But this research is not that well linked to the offline nature of the users, which till today remains a research gap. Moreover, the nature of Twitter discussions has been also described as ad hoc, which implies that the findings for one discussion are not applicable to another one. But both research on Twitter gateway formation [18] and our own research shows that technical nature of Twitter discussions has stable features that enable comparative research.

And, perhaps most important, several studies of Twitter have linked user rankings and influencer status to their offline metadata. Thus, studies on Twitter segments that belong to the USA and Sweden have shown that experts, professionals, and old-established organizations exert more considerable influence within discussions [26, 27]. In case of emergency, e.g. natural disaster or a terror attack, the so-called utility accounts may start to play a bigger role [24]. The same author has also shown that media and politicians belong to influencers on Twitter [24]; other works have shown that, for example, Norwegian politicians and journalists share patterns of Twitter activity [28] but not in all features. E.g. politicians posted significantly more retweets with hashtags than the journalists, and the politicians' promoted hashtags through retweets, while journalists used hashtags to promote original content [28].

Thus, among the various factors that are, arguably, expected to shape growth of a user's influencer status, beyond the platform and network nature, the most relevant for institutionalized users would be their activity, their belonging to certain types of institutions, their grassroots/systemic status, and their belonging to the majority/minority. But this has not yet been tested in comparative perspective, including the democratic/non-democratic states.

As it is known from a lot of research on the aftermath of the Arab Spring [29, 30], as well as from other studies on Twitter and authoritarianism [30, 31], post-communist, semi-authoritarian and authoritarian states may have differing patterns of political use of Twitter, as well as regime responses to it. This creates a perspective for comparing online discussion patterns in democratic, non-democratic, and hybrid or semi-authoritarian countries, to learn how much they actually differ.

2.2 Measuring Influence in Twitter Discussions

To dig a bit deeper into the nature of linkages between the offline user status and his/her online influencer status, we need to look at two major research lines that deal with: (1) how influencers are defined; and (2) how online behavior patterns of political actors is linked to their reaching the influencer positions.

As noted above, hierarchization of Twitter users has been documented in many studies, but not that many of them did it in relation to political actors. Moreover, as soon as it concerns political topics on Twitter, the overwhelming majority of the studies include into 'politically relevant actors' not only political institutions and parties, individual politicians, and grassroots movements and their leaders, but also media of various kind. We will follow this tradition and further discuss how offline dominance of these communicative actors translates into their online influencer status.

Defining Influencers. As we have addressed this issue in our previous research papers [7, 8], we will now focus on the aspects relevant for the topic of this research only, as well as on the newest literature in the field.

As stated previously, today's research on Twitter influencers may be divided in two streams. The first group of scholars includes an account or a group of users into influencers on the basis of non-network-dependent metrics such as the number of followers, comments, or retweets. Within this stream of literature, retweets have been marked as the preferred measure for defining who the influencers are – these users are retweeted both most frequently and most widely [32–34]. 'An analysis of responses and retweets provides a useful indication of the overall visibility of each account: as discussed, retweets are a means of amplifying the reach of a tweet, and thus of increasing the visibility of a tweet and its sender' [14, e.g. 17, 35, 36]. There even is a classification introduced by authors [37] for the most active users that includes politicians, bloggers and journalists as well as anonymous users and is based on two types of retweeting – that is, mentions via @ and RT networks.

Diffusion of messages on Twitter has different velocity depending on retweet activity [18] and user status [37, 38]. What is also important is that a retweet is considered to be 'a form of further disseminating someone else's message to one's own network' [4, 39]. Measuring number of retweets is about 'both influence in diffusion and value in market of a tweet' [40: p. 792]. A simple retweet is, in most cases, regarded as the sign of solidarization with the position expressed in the original tweet, but we also know that the retweeting practices vary considerably [35]. 'Retweeting users may even see themselves as information brokers, bridging distinct communities of interest by passing on tweets from one network cluster to another' [14]. In other works, comments and likes, as well as their combinations, are also considered important as influencer markers, but this line of research is definitely smaller.

The second approach measures connectivity (involvement into discussion or in a network) instead of activity. According to this approach, the most influential are those users who have potential in uniting echo chambers in Twitter because they form big nodes in the discussion networks. Most often, to measure this user potential, the metrics of SNA are employed. The clear difference between two approaches is rooted in fact that such SNA metrics as betweenness and pagerank centralities are network-dependent. "Degree centrality is conceptually the simplest one, which is defined as the number of links incident upon a node. Meanwhile, betweenness centrality and closeness centrality are global metrics which more complex but it can better to identify the influential nodes" [41]. A lot of combinations of network measures have been offered by researchers to define the top users; we rely most on (in/out)degree, betweenness, and pagerank metrics for ad hoc discussions [7].

In our previous research, we have pioneered juxtaposing absolute metrics and network-based metrics, to show that these correlate [7, 8]. In the most recent studies, there have also been attempts to correlate individual indicators, but no more than three [e.g. 4]. But this time we aim at linking this research method to the offline user status that we define via user metadata and user descriptions, as well as the content of the accounts.

User behavior and their influencer status. Several studies [28, 42, 43] have studied peculiar behavior that distinguishes them from average user profiles. Thus, it was shown that politicians and journalists form isolated networks and tend to interact within these circles – which also may be caused by the nature of the Twitter networks themselves, as other studies have found it for other types of users [7, 8]. Thus, more than two thirds of German politicians that actively use Twitter referred to other politicians in their tweets [5]. On the other hand, for wider audiences, 'official' accounts are able to establish themselves as authoritative sources of information, even in the open environments of social media' [14]. We also know that behavior matters: witty and sharply formulated political tweets, for example, have a great potential of spreading around the network [44, 45]. Also, candidates who have a prominent position in the media are generally also the ones who are more popular on Twitter [44].

Similar patterns may be found in online behavior of media and journalists. In case of Ferguson, despite having vast following, on average more than a million followers, legacy media journalists and organizations reciprocated by following only small numbers of users and were the least engaging in terms of retweeting and commenting other users [4]. These signs of traditional communicative hierarchies let journalists and legacy media remain within top popular users as defined via following but in several days were no longer the most influential in terms of generating discussions and gate-keeping information about the Ferguson decision. This is partly explained by their passive activity on Twitter, according to the authors.

Thus, we will try to explore whether patterns of political actors (including media) are similar in cross-cultural perspective.

2.3 Establishing Comparative Perspective: Twitter and Politics in Germany, the USA, and Russia

To form better our hypotheses, we need to establish a comparative perspective. For this, we will describe the political use of Twitter in the three countries and show that the cases under our scrutiny are comparable.

Twitter and politics in Germany, the USA, and Russia. Although Twitter technologically is the same all over the world, societal structure of users, as well as political clusterization and political use of Twitter, differs from country to country.

24% of US population used Twitter in 2016 (Pew Research Center). It is hard to identify the leaders of the US Twitter for the American audience, as the data that are provided are not country-based; e.g. the third poplar profile belongs to Barack Obama and is followed by 89.5 Mln of users who obviously live all over the world. According to the open data, most popular accounts of US citizens belong to the show business celebrities.

Throughout the recent years, online polarization has been increasing on the US Twitter [46, 47]. Support for the Democratic party is more visible on Twitter than that of Republicans (31% vs. 19%) [48]. Almost 60% of social media users feel irritated while participating in online discussions with users of different political position [49]. A very small number of large organizations that have a very large online network of Twitter followers are concentrating the online political voice [50]. Among members of the US House of Representatives, politicians with extreme ideological positions tend to have more Twitter followers, both left- and right-wing extremists have larger Twitter readership than their moderate peers [51]. Twitter is actively used in the USA in the election periods, as well as in times of nationwide events.

The divergence of users based on race is more significant in the US Twitter than elsewhere. Also, studies showed that education and ethnic heterogeneity are associated with uncivil communication on the US Twitter, while annual household income and unemployment rate were not related to Twitter incivility [52].

Russian Twitter was launched in 2011 and is less popular than the national social networks Vkontakte (desktop monthly reach in June 2016 46 Mln) and Odnoklassniki (desktop monthly reach in June 2016 30 Mln) [53] as well as Facebook (monthly reach in June 2016 14.4 Mln) [54]. Monthly reach of Twitter was about 12 Mln [54, 55]. Kelly et al. have revealed several political clusters within the Russian Twitter in 2010 and 2011. The cleavages between clusters correlated with topics of discussion rather than with social stratification, which is a sign of presence of echo chambers in Russian Twitter [56]. At the same time, both pro-government and oppositional political voices are represented in Russian Twitter almost equally, with slight dominance of pro-establishment voices [57, 58]. Leaders of the Russian Twitter are TV anchors and politicians: prime minister Dmitry Medvedev and TV showman Ivan Urgant have about 5.5 Mln followers each. Old-established news organizations (*Channel One* and *Vesti*) are followed by 3.7 Mln of Twitter users each [59]. Parody microblogs are also highly visible: the spoof account @KermlinRussia gained over 1.5 Mln followers becoming more popular than the official presidential account the spoof was parodying [60]. While not a big share of Russian population uses Twitter the platform attracts more politically active users [22, 61].

In general, the Russian segment of Twitter is under-researched. The evidence from previous research provides contradictory data on whether Russian Twitter tends to form echo chambers [56] or play rather a role of mediator between different views [22, 61, 62]. Also, very recently, the number of bots has risen substantially, especially in discussions on political matters, but this was still irrelevant in 2013.

Twitter is also not very popular in Germany. While more than 30% of German population have created once a Twitter account, in 2015 only about 10% of German population used this micro-blogging platform actively [63]. Leaders of the German Twitter are celebrities - football and show business stars with about 4 Mln of followers. In comparison to them, most popular media accounts gained about 2 Mln of followers and accounts of politicians attracted less than 400,000 followers.

Among German Twitter users who were politically active in the 2013 elections, almost every third prefers to participate in supportive networks where users always or often agree with others [64]; thus, German Twitter is tending to form echo chambers and user chains. Political journalists unite in a journalism-centered bubble where political accounts are included since journalists prefer to follow institutionalized users as sources and do not use Twitter for communication with the audience [65].

This description shows that we may expect higher presence of political actors in Twitter of democratic countries, and these users will be of different nature (institutional as well as grassroots, perhaps more institutional in Germany than in the US). We can also expect presence of all major political powers on Twitter, but their interaction with users will be scarce, as previous research shows.

The conflicts under scrutiny. The three conflicts that we have chosen for our analysis include the Biryulyovo bashings of 2013, Russia, the Ferguson unrest of 2014, USA, and the Cologne mass harassment of 2016, Germany. These conflicts were selected due to the fact that they contained clear political elements, were all ethnic/racial-based, provoked massive discussions on Twitter (reaching the national Twitter trending topics), and had similar ad hoc nature in terms of technical organization. Political criteria for selection included the following: the cases had a violent trigger, evoked or revealed major political polarization in the society, caused peaceful protest, and involved institutional political actors and police forces.

In Russia, major anti-migrant bashings happened in September 2013 in the Moscow district of Biryulevo, triggered by an alleged killing of a Muscovite Egor Sviridov by an Uzbek Orkhan Zeinalov and directed to a warehouse where a lot of Central Asian immigrants (mostly illegal) lived and traded. Police and additional military forces had to gather in the area to stop violence; there were claims on the part of attackers that police over-reacted to the events. Later that autumn, peaceful gatherings of local dwellers continued; there were no direct links in change of Moscow policy towards immigration, but the markets and warehouses in the Russian capital have since then been under particular scrutiny of police and immigration officers.

In the USA, an African American teenager Mike Brown was killed unarmed by a white police officer named Darren Wilson. This event provoked major riots not only in the town of Ferguson where the killing happened but also in other cities and towns; peaceful protest spread up to Washington, DC. The case has polarized the audiences; several media reported that more money was raised for Wilson than for the Brown family. For several weeks, protests continued and had minor policy implications on the local level.

In Germany, in the New Year eve of 2016, mass harassment of female celebrators happened on the main square of Cologne; over a 1,000 women self-reported to be affected, and several risked their lives due to firework pieces being directed to them. As police proved, most of the harassers belonged to the re-settled communities from Middle East and North African countries. The conflict provoked nationwide debates but was also characterized by the silence of media for the first few days, as editorial guidelines did not let papers and TV report on the nationality of harassers.

3 Method

3.1 The Research Questions and the Hypotheses

Based on what was said above, we formed the following research questions:

RQ1. Is the structure of presence of politically relevant (institutionalized political, grassroots political, and media) users among the influencers similar in all the three cases?

RQ2. Do politically institutionalized users play bigger role in the discussion than other politically aligned users?
RQ3. Is offline status of the users linked to their online influencer status, and by which parameters?

According to the research questions, we have formulated three hypotheses:

H1. Due to the nature of the civil society in the countries, there is bigger presence of grassroots political users and NGOs in the USA and Germany in the top user lists than in Russia.
H2. Due to importance of political actors during conflicts, as well as inertia of traditional hierarchies, users who belong to the current national-scale political institutions form big nodes in the discussion networks – that is, they are ranked high by degree, betweenness, and pagerank centralities.
H3. The position of users in the top lists correlates with the users' political status. Thus, the more politically institutionalized a user is, the more he is 'liked', retweeted, and commented; also, (s)he occupies a higher place in the structural hierarchy of the discussion web graph (higher InDegree, Degree, Betweenness, and Pagerank). As to the OutDegree, it inversely correlates with political institutionalization, as traditional hierarchies are reproduced in the online discussions.

3.2 The Research Methodology

To answer the research questions, we have done the following:

- We conducted vocabulary-based web crawling to collect the discussion content and to analyze their structure. The keywords were selected at www.trendinalia.com, in accordance with hashtags high-rated by this web service that monitor Twitter trending topics around the world in real time environment and diachronically. A special web crawler has been developed to bypass the limitations of Twitter API;
- We have reconstructed the discussion web graphs to get the data on network-related metrics;
- We sampled 50 top users based on nine parameters (number of tweets; number of interactions (likes, comments and retweets); centralities – degree, indegree, out-degree, betweenness, and pagerank); then we merged them in general lists of top users. Thus, if a user was found in at least one single-metric list, he/she was to be found in the merged top user list. As many users were within top lists by many parameters, the final top lists included 205 users for Russia, 205 users for Germany, and 230 users for the USA. The total number of users in the datasets included 3574 users for Russia, 12382 users for Germany, and 70018 users for the USA;
- We manually coded the users in the lists for their offline status, aiming at showing their growing political institutionalization (from 'ordinary users' to 'institutional political actors');
- We used descriptive statistics (Spearman's rho and Cramer's V) to see whether the offline user status correlates with the online metrics;
- We qualitatively assessed the place of political users within the top lists.

4 Findings

H1. Due to the nature of the civil society in the countries, there is bigger presence of grassroots political users and NGOs in the USA and Germany in the top user lists than in Russia (Table 1).

Table 1. User status in the top user lists

	Germany N = 230	Russia N = 205	USA N = 205
Ordinary people	48.70%	36.59%	27.32%
Activists and politically active bloggers	10.87%	13.66%	23.41%
Media	23.48%	22.93%	17.56%
Political actors outside state institutions	1.74%	3.41%	3.41%
State authorities of any level and their individual representatives	3.48%	6.83%	0.49%
Arts people, celebrities, and experts	6.96%	3.41%	13.66%
Other/non-defined/irrelevant	4.78%	13.17%	14.15%

The biggest share of grassroots political users is present in the USA – almost 25% of our sample in comparison to 11% in Germany and 14% in Russia. In Germany, many ordinary users have a clear political position that can be understood from their tweets but they don't define themselves as activists. Despite the share of authorities is bigger in Russia than in Germany, the level of authorities should be taken into consideration. While the sample of Russian political users includes low-level politicians - deputies of local councils or executives, as well as institutionalized party members outside the State Duma, the list of influential German Twitter users includes a member of the European Parliament, former Federal Minister of Justice and Consumer Protection, and members of parliamentary parties. 50% of politicians belonging to the German Twitter influencers represent the right-wing party Alternative for Germany (AfD). Thus, H1 is partly supported for the US and Russia.

H2. Due to importance of political actors during conflicts, as well as inertia of traditional hierarchies, users who belong to the current national-scale political institutions form big nodes in the discussion networks – that is, they are ranked high by degree, betweenness, and pagerank centralities.

H2 is not supported for the United States, since there are no state authorities of any level and their individual representatives among users ranked high by degree, betweenness, or pagerank centralities. For Germany, it is partly supported, since there are representatives of 'Alternative for Germany' (rating 11 out of 50 by betweeness and 44 by pagerank), 'Die Linke' (The Left Party – rating 33 by degree) and the Social Democratic Party of Germany (rating 41 for degree).

Representatives of state actors are much more authoritative on the Russian Twitter. There are 7 of them that were ranked high by degree (50% of this type of users in the whole sample), 4 ranked high by betweenness and 8 ranked high by pagerank.

Among these users, one is the member of Committee on Drug Prevention by the Federation Council, upper house of the Federal Assembly of Russia (6 by degree, 10 by betweenness, 15 by pagerank). Beside him, the three users also were ranked high by all the three metrics: a member of the Council of the youth wing of the 'United Russia' party, an account of the Civic Chamber of the Russian Federation and a member of a populist Liberal Democratic Party of Russia.

H3. The position of users in the top lists correlates with the users' political status. Thus, the more politically institutionalized a user is, the more he is 'liked', retweeted, and commented; also, (s)he occupies a higher place in the structural hierarchy of the discussion web graph (higher InDegree, Degree, Betweenness, and Pagerank). As to the OutDegree, it inversely correlates with political institutionalization, as traditional hierarchies are reproduced in the online discussions.

What we see from Table 2 is the following. First, for Germany and the USA, we see a similar pattern. The growth of the level of political institutionalization correlates with absolute figures of user activity, namely likes and retweets – that is, political figures tend to maintain authority and keep the audience's wish to interact with their content. This is a sign of traditional hierarchies still being in place in online discussions; belonging to an institution linked to politics slightly raises the chances to be more times liked and retweeted – but not commented, which also shows that the nature of the discussion is not dialogue and the politically institutionalized users do not tend to enter discussions and generate mutual commenting. This is also partially supported by the fact that there is no correlation between user status and betweenness centrality – that is, politically institutionalized users do not tend to be linking groups of users more than others, which they could do if talked to people more.

Table 2. Spearman's rho correlating the user status to…

	…likes	…retweets	…comments	…InDegree	…OutDegree	…Degree	…Betweenness	…Pagerank
USA	0.229**	0.235**	–	0.214**	−0.179*	–	–	–
Germany	0.145*	0.231**	–	0.150*	−0.227**	–	–	0.189**
Russia	–	–	–	0.257**	–	0.231**	0.153*	0.288**

Also, for the USA and Germany, we see clear support of the idea that political institutionalization raises the number of users who interact with the given user (by InDegree metric) and diminishes the number of users who are interacted with on the part of the given user. In other words, if you belong to a political institution, you post information that people tend to like and share, and you remain attractive for more users than average, but you tend to be passive in communicating with them. This is another sign of traditional hierarchies rising in the ad hoc discussions where all the users are seemingly equal. In Germany, political users also show higher authority among the most authoritative users (by pagerank centrality), but we explain this by the nature of the case and the top user list – while in Ferguson there were almost no politicians at all, in Germany their presence was bigger and they could enter the 'authority chain' of users with high pagerank centrality.

In Russia, by sharp contrast, we see that political users, despite their relatively massive presence, do not create the picture of being different from ordinary users in terms of user attention (likes, retweets, or comments), but the structure of the discussion still makes them different from the average user. They are not only referenced to by more people (not because of content, it seems, but because of their status), but the people who reference to political actors also tend to be more authoritative. It even seems that institutionalized users form something like a self-referential cluster of the discussion, with external audience seeing also referencing to it. This looks like another form of keeping hierarchies – based not on political individual but on political status itself.

But anyway, in all the three countries, politically institutionalized users do not stand in the position of linking people in the discussions and forming circles of interest around them – only in Russia their betweenness centrality is slightly connected to the user status but the correlation is really weak. This may be the sign of the situation when, while traditional hierarchies still work in people's minds, the real role of political actors in public discussions is not that significant.

These findings are partly supported also by Cramer's V metric (see Table 3) – with the difference for Germany where there are practically no significant correlations that would show up. But for the USA and Russia correlations not only are there but are really strong, showing the differences between content-based and status-based nature of influencing the discussion by politically institutionalized actors in the US and Russia, respectively.

Table 3. Cramer's V correlating the user status to…

	…likes	…retweets	…comments	…InDegree	…OutDegree	…Degree	…Betweenness	…Pagerank
USA	0.925**	0.944**	0.915**	–	–	–	–	–
Germany	–	–	–	0.813 *(sig. 0.055)*	–	–	–	–
Russia	–	–	–	0.716***	–	0.747***	–	0.980**

Thus, we see our hypothesis is partly supported for Germany and the USA; there is the pattern of attention growth for more politically institutionalized users there, based on the number of users who pay attention to politicians, as well as on their content which users find sharable in the crisis situations. In Russia, on the contrary, the pattern is not there, and the discussion first seems more deprived of hierarchies, but the structural metrics reveal that the status itself matters – arguably, more than content these user post. This leads us to the conclusion that political presence in Twitter in democratic countries differs from that in post-authoritarian ones, especially in terms of what creates the authority.

5 Conclusion

Thus, we see that even preliminary research like ours shows that there are significant gaps in political representation within conflictual discussions on Twitter; there is room for significant improvement in terms of discussion engagement of political actors, including media and grassroots politicians.

We also see that the USA and Germany tend to differ from Russia in terms of how exactly the hierarchies are reproduced – in Russia, it is the status itself that matters, and the more politically institutionalized a user is, the bigger chance is that he/she will form an important network node, while in the US and Germany we see that content matters more, but political users fail anyway to be discussion centers.

This provides new grounds for further research on behavior and role of politically aligned users on Twitter.

Of course we see many limitations in this kind of research; also, we have not provided in-depth analysis of the strategies of political actors and the causes that lead them to the detected (and relative) failure in uniting the user groups and discussion echo chambers. But this lies a bit beyond the scope of this paper and demands further research. What we would like to underline is that the divergent patterns of political influencing may be grounded not only in the nature of the national publics but also the nature of the national political arenas.

Acknowledgements. This research has been supported in full by Russian Science Foundation, research grant 16-18-10125.

References

1. Bennett, W.L., Segerberg, A.: The Logic of Connective Action: Digital Media and the Personalization of Contentious Politics. Cambridge University Press, Cambridge (2013)
2. Bastos, M.T., Mercea, D.: Serial activists: political Twitter beyond influentials and the twittertariat. New Media Soc. **18**(10), 2359–2378 (2015). doi:10.1177/1461444815584764
3. Bonilla, Y., Rosa, J.: #Ferguson: digital protest, hashtag ethnography, and the racial politics of social media in the United States. Am. Ethnol. **42**, 4–17 (2015). doi:10.1111/amet.12112
4. Groshek, J., Tandoc, E.: The affordance effect: gatekeeping and (non)reciprocal journalism on Twitter. Comput. Hum. Behav. **66**, 201–210 (2017)
5. Baumann, A., Fabian, B., Lessmann, S., Holzberg, L.: Twitter and the political landscape – a graph analysis of German politicians. Research Papers, 132 (2016)
6. Lasorsa, D.L., Lewis, S.C., Holton, A.E.: Normalizing Twitter. Journal. Stud. **13**, 19–36 (2012). doi:10.1080/1461670X.2011.571825
7. Bodrunova, S.S., Litvinenko, A.A., Nigmatullina, K.R., Blekanov, I.S., Smolarova, A.S.: Ad-hoc publics and influencers on the Russian Twitter: the case of discussion on migrants. In: Proceedings of 3rd International Multidisciplinary Scientific Conferences on Social Science and Arts (SGEM - 2016). Volume II: Sociology and Healthcare, Albena, Bulgaria, pp. 73–80, 22–31 August 2016
8. Bodrunova, S.S., Litvinenko, A.A., Blekanov, I.S. Influencers on the Russian Twitter: institutions vs. people in the discussion on migrants. In: Proceedings of ACM International Conference 'Electronic Governance and Open Society: Challenges in Eurasia', St. Petersburg, Russia, pp. 212–222, 21–23 November 2016
9. Toepfl, F.: Managing Public Outrage: Power, Scandal, and New Media in Contemporary Russia. In: Proceedings of the Conference 'Internet, Politics, Policy 2010: An Impact Assessment', Oxford, UK (2010). http://microsites.oii.ox.ac.uk/ipp2010/programme/79
10. Sherstobitov, A.: The potential of social media in Russia: from political mobilization to civic engagement. In: ACM International Conference Proceeding Series, pp. 162–166 (2014). doi:10.1145/2729104.2729118

11. Smyth, R., Oates, S.: Mind the gaps: media use and mass action in Russia. Europe-Asia Stud. **67**(2), 285–305 (2015). doi:10.1080/09668136.2014.1002682
12. Denisova, A.: Democracy, protest and public sphere in Russia after the 2011-12 anti-government protests: digital media at stake. Media Cult. Soc. (2016). doi:10.1177/0163443716682075
13. Bodrunova, S., Litvinenko, A.: Fragmentation of society and media hybridisation in today's Russia: how Facebook voices collective demands. Zhurnal Issledovanii Sotsial'Noi Politiki **14**(1), 113–124 (2016)
14. Bruns, A., Burgess, J.: Researching news discussion on Twitter. Journal. Stud. **13**, 801–814 (2012). doi:10.1080/1461670X.2012.664428
15. Bakshy, E., Hofman, J.M., Mason, W.A., Watts, D.J.: Everyone's an influencer: quantifying influence on Twitter. In: Proceedings of International Conference on Web Search and Data Mining (WSDM), pp. 65–74. ACM, Hong Kong (2011)
16. Habermas, J.: Political communication in media society: does democracy still enjoy an epistemic dimension? The impact of normative theory on empirical research. Commun. Theory **16**(4), 411–426 (2006). doi:10.1111/j.1468-2885.2006.00280.x
17. Kwak, H., Lee, C., Park, H., Moon, S.: What is Twitter, a social network or a news media? In: Proceedings of the 19th International Conference on World Wide Web, Raleigh, North Carolina, USA (2010). doi:10.1145/1772690.1772751
18. Bastos, M.T., Raimundo, R.L.G., Travitzki, R.: Gatekeeping Twitter: message diffusion in political hashtags. Media Cult. Soc. **35**(2), 260–270 (2013). doi:10.1177/0163443712467594
19. Wu, S., Hofman, J.M., Mason, W.A., Watts, D.J.: Who says what to whom on Twitter. In: Proceedings of the 20th International Conference on World Wide Web, pp. 705–714. ACM, Hyderabad (2011)
20. Juris, J.: Reflections on #OccupyEverywhere: social media, public space and emerging logics of aggregation. Am. Ethnol. **39**, 259–279 (2012)
21. Clark, L.S.: Constructing public spaces| participants on the margins: examining the role that shared artifacts of engagement in the Ferguson protests played among minoritized political newcomers on Snapchat, Facebook, and Twitter. Int. J. Commun. **10**, 235–253 (2016)
22. Bodrunova, S.S., Litvinenko, A.A.: New media and the political protest: the formation of a public counter-sphere in Russia of 2008-12. In: Makarychev, A., Mommen, A. (eds.) Russia's Changing Economis and Political Regimes: The Putin Years and Afterwards, pp. 29–65. Routledge, London (2013)
23. Dahlgren, P.: Media and Political Engagement: Citizens, Communication and Democracy. Cambridge University Press, Cambridge (2009)
24. Bruns, A.: Gatekeeping, gatewatching, real-time feedback: new challenges for journalism. Braz. Journal. Res. J. **7**(2), 117–136 (2011)
25. Bruns, A.: Gatewatching: Collaborative Online News Production. Peter Lang, New York (2005)
26. Fox, S., Zickuhr, K., Smith, A.: Twitter and status updating, fall 2009. Pew Internet and American Life Project. http://www.pewinternet.org/2009/10/21/twitter-and-status-updating-fall-2009. Accessed 22 May 2017
27. Ruth, P.: The linguistics of self-branding and micro-celebrity in Twitter: the role of hashtags. Discourse Commun. **6**(2), 181–201 (2012). doi:10.1177/1750481312437441
28. Enli, G., Simonsen, C.-A.: 'Social media logic' meets professional norms: Twitter hashtags usage by journalists and politicians. Inf. Commun. Soc. **20**, 1–16 (2017). doi:10.1080/1369118X.2017.1301515
29. Lynch, M.: After Egypt: the limits and promise of online challenges to the authoritarian Arab state. Perspect. Polit. **9**, 301–310 (2011)

30. Tufekci, Z., Wilson, C.: Social media and the decision to participate in political protest: observations from Tahrir Square. J. Commun. **62**(2), 363–379 (2012). doi:10.1111/j.1460-2466.2012.01629.x
31. Howard, P.N., Agarwal, S.D., Hussain, M.M.: When do states disconnect their digital networks? Regime responses to the political uses of social media. Commun. Rev. **14**(3), 216–232 (2011). doi:10.2139/ssrn.1907191
32. Freberg, K., Graham, K., McGaughey, K., Freberg, L.A.: Who are the social media influencers? A study of public perceptions of personality. Publ. Relat. Rev. **37**, 90–92 (2011)
33. Aquino, J.: Boost Brand Advocates and Social Media Influencers (2013). http://www.destinationcrm.com/Articles/Editorial/Magazine-Features/Boost-Brand-Advocates-and-Social-Media-Influencers-86675.aspx. Accessed 19 May 2017
34. Sajuria, J., van Heerde-Hudson, J., Hudson, D., Dasandi, N., Theocharis, Y.: Tweeting alone? An analysis of bridging and bonding social capital in online networks. Am. Polit. Res. (2015). doi:10.1177/1532673X14557942
35. Boyd, D., Golder, S., Lotan, G.: Tweet, tweet, retweet: conversational aspects of retweeting on twitter. In: Proceedings of 43rd Hawaii International Conference on System Sciences (HICSS), pp. 1–10. IEEE, Kauai (2010)
36. Cha, M., Haddadi, H., Benevenuto, F., Gummad, K.P.: Measuring user influence on twitter: the million follower fallacy. In: Proceedings of the 4th International AAAI Conference on Weblogs and Social Media, Washington, D.C., pp. 10–17 (2010)
37. Larsson, A.O., Moe, H.: Studying political microblogging: Twitter users in the 2010 Swedish election campaign. New Media Soc. **14**, 729–747 (2012). doi:10.1177/1461444811422894
38. Colliander, J., Marder, B., Falkman, L.L., Madestam, J., Modig, E., Sagfossen, S.: The social media balancing act: testing the use of a balanced self-presentation strategy for politicians using Twitter. Comput. Hum. Behav. **74**, 277–285 (2017). doi:10.1016/j.chb.2017.04.042
39. Mascaro, C., Goggins, S.: Twitter as virtual town square: citizen engagement during a nationally televised Republican primary debate. In: American Political Science Conference, New Orlean, LA (2012)
40. Liu, G., Shi, C., Chen, Q., Wu, B., Qi, J.: A two-phase model for retweet number prediction. In: Li, F., Li, G., Hwang, S-w, Yao, B., Zhang, Z. (eds.) WAIM 2014. LNCS, vol. 8485, pp. 781–792. Springer, Cham (2014). doi:10.1007/978-3-319-08010-9_84
41. Maharani, W.: Identifying influential user in Twitter: analysis of tweet content similarity in weighted network. J. Theor. Appl. Inf. Technol. **81**(3), 558–563 (2015)
42. Graham, T., Broersma, M., Hazelhoff, K., van't Haar, G.: Between broadcasting political messages and interacting with voters. Inf. Commun. Soc. **16**, 692–716 (2013). doi:10.1080/1369118X.2013.785581
43. Larsson, A.O., Ihlen, Ø.: Birds of a feather flock together? Party leaders on Twitter during the 2013 Norwegian elections. Eur. J. Commun. **30**, 666–681 (2015). doi:10.1177/0267323115595525
44. Aelst, P., van Erkel, P., D'heer, E., Harder, R.A.: Who is leading the campaign charts? Comparing individual popularity on old and new media. Inf. Commun. Soc. **20**(5), 715–732 (2017). doi:10.1080/1369118X.2016.1203973
45. Parmelee, J.H., Bichard, S.L.: Politics and the Twitter Revolution: How Tweets Influence the relationship Between Political Leaders and the Public. Lexington Books, Lanham (2012)
46. Garimella, V.R.K., Weber, I.: A Long-Term Analysis of Polarization on Twitter. Pre-print of a Short Paper Accepted at ICWSM 2017. arXiv:1703.02769v2
47. Cohen, R., Ruths, D.: Classifying political orientation on Twitter: it's not easy! In: Proceedings of the Seventh International AAAI Conference on Weblogs and Social Media (2013)

48. Greenwood, S., Perrin, A., Duggan, M.: Social Media Update 2016 (2016). http://www.pewinternet.org/2016/11/11/social-media-update-2016/. Accessed 19 May 2017
49. Duggan, M., Smith, A.: The Political Environment on Social Media (2016). http://www.pewinternet.org/2016/10/25/the-political-environment-on-social-media/. Accessed 19 May 2017
50. Hong, S., Nadler, D.: Social media and political voices of organized interest groups: a descriptive analysis. In: Proceedings of the 16th Annual International Conference on Digital Government Research, pp. 210–216. ACM, Phoenix (2015)
51. Hong, S., Kim, S.H.: Political polarization on Twitter: implications for the use of social media in digital governments. Gov. Inf. Q. **33**, 777–782 (2016). doi:10.1016/j.giq.2016.04.007
52. Vargo, C.J., Hopp, T.: Socioeconomic status, social capital, and partisan polarity as predictors of political incivility on Twitter: a congressional district-level analysis. Soc. Sci. Comput. Rev. **35**(1), 10–32 (2017). doi:10.1177/0894439315602858
53. TNS Web Index 2016, June, 2016 http://mediascope.net/. Accessed 19 May 2017
54. Frolova, E.: The most popular social networks in Russia-2016 (2016). pro-smm.com. Accessed 19 May 2017
55. Brand Analytics. Social Networks in Russia (2016). blog.br-analytics.ru. Accessed 19 May 2017
56. Kelly, J., Barash, V., Alexanyan, K., Etling, B., Faris, R., Gasser, U., Palfrey, J.: Mapping Russian Twitter. Berkman Center for Internet and Society, Harvard University, Cambridge (2012)
57. Greene, S: Twitter and the Russian Street: Memes, Networks and Mobilization. CNMS Working Paper 2012/1. Center for the Study of New Media and Society, New Economic School, Moscow (2012)
58. Nikiporetz-Takigawa, G.: Tweeting the Russian protests. Digit. Icons: Stud. Russ. Eurasian Cent. Eur. New Media **9**, 1–25 (2013)
59. SocialBakers.com. Twitter Statistics for Russia. https://www.socialbakers.com/statistics/twitter/profiles/russia. Accessed 19 May 2017
60. Denisova, A.: Parody microbloggers as chroniclers and commentators on Russian political reality. Demokratizatsiya: J. Post-Sov. Democratization **25**(1), 23–41 (2017)
61. Bodrunova, S.S., Litvinenko, A.A.: Four Russias in communication: fragmentation of the Russian public sphere in the 2010s. In: Dobek-Ostrowska, B., Glowacki, M. (eds.) Democracy and Media in Central and Eastern Europe 25 Years On, Wroclaw, pp. 63–79 (2015)
62. Barash, V., Kelly, J.: Salience vs. commitment: dynamics of political hashtags in Russian Twitter. Berkman Center Research Publication (2012)
63. Mosca, L., Quaranta, M.: News diets, social media use and non-institutional participation in three communication ecologies: comparing Germany, Italy and the UK. Inf. Commun. Soc. **19**(3), 325–345 (2016)
64. Vaccari, C., Valeriani, A., Barbera, P., Jost, J.T., Nagler, J., Tucker, J.A.: Of echo chambers and contrarian clubs: exposure to political disagreement among German and Italian users of Twitter. Soc. Media+ Soc. **2**(3), 1–24 (2016). doi:10.1177/2056305116664221
65. Nuernbergk, C.: Political journalists' interaction networks: the German federal press conference on Twitter. Journal. Pract. **10**(7), 868–879 (2016)

The STEP Project: Societal and Political Engagement of Young People in Environmental Issues

Maria Vogiatzi[1(✉)], Christodoulos Keratidis[1], Manos Schinas[2], Sotiris Diplaris[2], Serdar Yümlü[3], Paula Forbes[4], Symeon Papadopoulos[2], Panagiota Syropoulou[1], Lazaros Apostolidis[2], Ioannis Kompatsiaris[2], and Machi Symeonidou[1]

[1] DRAXIS Environmental SA, Mitropoleos 63, 54623 Thessaloniki, Greece
mvogiatzi@draxis.gr
[2] Centre for Research and Technology Hellas, Information Technologies Institute, 6th Km. Charilaou-Thermi Road, Thermi, 57001 Thessaloniki, Greece
[3] Sampas Bilisim Ve Iletisim Sistemleri Sanayi Ve Ticaret A.S., Çubuklu Mahallesi Orhan Veli Kanık Caddesi Eryılmaz, 34810 Istanbul, Turkey
[4] Abertay University, Bell Street, Dundee, Scotland

Abstract. Decisions on environmental topics taken today are going to have long-term consequences that will affect future generations. Young people will have to live with the consequences of these decisions and undertake special responsibilities. Moreover, as tomorrow's decision makers, they themselves should learn how to negotiate and debate issues before final decisions are made. Therefore, any participation they can have in environmental decision making processes will prove essential in developing a sustainable future for the community.

However, recent data indicate that the young distance themselves from community affairs, mainly because the procedures involved are 'wooden', politicians' discourse alienates the young and the whole experience is too formalized to them. Authorities are aware of this fact and try to establish communication channels to ensure transparency and use a language that speaks to new generations of citizens. This is where STEP project comes in.

STEP (www.step4youth.eu) is a digital Platform (web/mobile) enabling youth Societal and Political e-Participation in decision-making procedures concerning environmental issues. STEP is enhanced with web/social media mining, gamification, machine translation, and visualisation features.

Six pilots in real contexts are being organised for the deployment of the STEP solution in 4 European Countries: Italy, Spain, Greece, and Turkey. Pilots are implemented with the direct participation of one regional authority, four municipalities, and one association of municipalities, and include decision-making procedures on significant environmental questions.

Keywords: E-participation · Social media mining · Environmental decision making · Youth engagement strategies

I. Kompatsiaris et al. (Eds.): INSCI 2017, LNCS 10673, pp. 148–156, 2017.
https://doi.org/10.1007/978-3-319-70284-1_12

1 e-Participation

Promoting youth participation is fundamental in the EU policy. It is an underlying theme of the EU Youth Strategy, and it is incorporated into the Treaty of Lisbon, where Article 165 TFEU stipulates that 'Union action shall be aimed at encouraging the participation of young people in democratic life in Europe'.

According to the findings of the Eurobarometer (Standard EB 77, spring 2012), half of the young people tend to distrust the European Union. Traditional channels of representative democracy, such as voting at elections and joining political parties only partially stimulate young people's interest in active participation. The percentage of young people voting in elections and being members of a political party is significantly lower than that of their elders [1].

As a reaction to this public discontent, many European countries have started exploring the potential of Information and Communication Technology (ICT) to regain citizens' trust and revitalise European democracy by developing a more responsive, transparent, and participatory decision-making process [2].

The participation of young people in decision making is especially relevant and extremely important in environmental issues. Decisions taken now on matters such as climate change, the depletion of resources, and the loss of biodiversity will have long-term consequences that will affect the future generations. Young people will have to live longer with the consequences of current decisions, and have special concerns and responsibilities in relation to the environment. Their participation in environmental decision making is an essential part of developing a sustainable future.

2 STEP Project

The STEP project (www.step4youth.eu) is supported by the European Union's Horizon 2020 research and innovation programme and aims to increase and support participation of young European citizens in decision making for environmental issues. STEP will develop and pilot test a cloud eParticipation platform that will facilitate interaction between policy makers and young people and will enable youth Societal and Political e-Participation in decision-making procedures concerning environmental issues. The main goal of the STEP project is to directly involve at least 8,000 participants in the respective environmental pilots and indirectly inform more than 80,000 people on e-participation.

The STEP Platform (www.step.green) is enhanced with a number of features, so as to become appealing to young people and engage them in e-participation processes. The platform was launched to the general public on February 2017 and is available as a mobile (both iOS and Android) and web application.

The main ideas underpinning the STEP project and the platform are briefly described below:

- e-participation features: The STEP platform offers a wide range of e-participation features that serve both top-down and bottom-up approaches.

- Integrated platform: STEP uses social media/web mining technologies that enable young people to view personalized and enriched information relevant to the topic they are viewing, filtered through multiple social media and web sources, in one unified environment.
- Engagement: Social media analytics and monitoring tools can be used by the authorities in order to facilitate the identification of the preferences of young people, and thus develop an effective strategy to engage young citizens.
- Removal of language barriers: Machine translation technology allows users to access content from other countries in their own language, removing language barriers and fostering cross-border interaction.
- Policy-making support: Visual representations of results that enable policy makers to interpret the results of their campaigns, view patterns and spot trends, facilitating decision making.

3 STEP Pilots

Six pilots in real contexts are being organised for the deployment of the STEP solution in 4 European Countries: Italy, Spain, Greece, and Turkey. Pilots are implemented with the direct participation of one regional authority, four municipalities, and one association of municipalities, and include decision-making procedures on significant environmental questions. During the STEP pilot, the public authorities invite young citizens to use the platform and participate in decision making processes on local environmental issues.

In order to assist the pilot partners for setting up their public participation processes effectively, a framework for public participation has been created, by combining information from available best practice guidelines and toolkits and adapting it to the STEP project specificities. This framework aims to provide a high-level consultation to the pilot authorities when deciding how best to execute the pilot needs in environmental decision-making.

Available best practice guidelines have been identified and reviewed, such as:

- Guidelines for successful e-participation by young people, Best practice manual [3]
- OPIN Guidelines – Draft created by EUth project [4]
- US Environmental Protection Agency Public Participation Guide [5]
- Public Participation in Government Decision-making, Better practice guide of Victorian Auditor-General's Office [6]
- Guidelines for successful e-participation by young people in decision-making at local, regional, national and European levels, International Youth Service of the Federal Republic of Germany [7].

These guidelines structure the public participation process by breaking it down to different phases, and present comprehensive recommendations for each phase in order to support the public authorities, or the operators to fully understand the procedures and implement in detail each phase of the public participation. Some of them are specifically related to e-participation of young people, while one focuses on environmental issues.

In addition to the guidelines above, the toolkits for public participation of the US Environmental Protection Agency in cooperation with the International Association for Public Participation (IAP2) and other organisations have been reviewed. These toolkits provide basic information and ideas as well as useful links to more content on the web. They follow a logical path to understand, plan, and implement a public participation process. Additionally, they are designed with government agencies in mind, to help those who must manage processes where public input is important to decision-making and provide a clear overview of important considerations in the design and implementation of a public participation initiative such as the STEP project's pilot. The following toolkits have been reviewed:

- Core Values for the Practice of Public Participation, International Association for Public Participation (IAP2) [8]
- Core Principles for Public Engagement, National Coalition for Dialogue & Deliberation (NCDD) [9]
- Participation and Civic Engagement, World Bank [10]
- Better Decisions through Consultation and Collaboration, US Environmental Protection Agency (EPA) [11]
- Spectrum of Public Participation, International Association for Public Participation (IAP2) [12]
- How to Consult and Involve the Public, US Environmental Protection Agency (EPA) [13]
- How to Evaluate Public Involvement, US Environmental Protection Agency (EPA) [14]

As every public participation procedure is unique, it should be planned based on its particularities. For this reason, the STEP public participation framework has been developed by selecting relevant elements from these guidelines and toolkits.

The STEP public participation includes 5 basic steps:

- STEP 1: Background: This step is about mapping the general context of the pilot, and starts at the beginning of the project.
- STEP 2: Planning: The planning step includes the main preparation activities and sets concrete goals, timeframe, responsibilities, and rules. This is the first step for realizing what has been sketched out in the Background.
- STEP 3: Action: The actual pilot execution of the e-participation process, once the environment is prepared and the whole process planned.
- STEP 4: Communication: This is a key step that includes all the communication activities which will be performed throughout the pilot lifetime. The most important mission of this step is to engage young people and increase STEP popularity.
- STEP 5: Feedback and evaluation: Feedback is an iterative process as the collected input by the platform users enhances and enriches the evaluation process, while the quality of the evaluation enables continuous improvement and learning through its implementation.

In order to define the local pilot plans in a structured and uniform way, a questionnaire was developed for the collection of information from the pilot partners. The questionnaire includes 10 sections. The questions and recommendations have been formulated by combining elements of the best practice guidelines, which have been selected according to their relevance to the project.

Through the questionnaire the public authorities are guided to define and explain the approach of conducting public participation and how this relates to the particular characteristics of the pilot and meet the needs of decision makers, public authority officers and young citizens. This process results in specifications for the public participation processes, and the input collected will be used to form the local pilot plans.

The questionnaire includes the following sections:

- Objectives and scope: Clear and comprehensive description of the objectives and scope of the participation procedure that the public authority intends to set up.
- Legal framework: Legislation that is relevant to public participation and regulations that the authority has to follow when making a decision.
- Selection of issues to be brought under public participation: Definition of specific issues on which young people will be consulted through the STEP platform.
- Public participation procedure: Definition of the public participation rules, how public participation will be conducted, and how the decision will be made.
- Timeframe: Definition of the steps and the time frame of the process in accordance with the general STEP pilot timeframe.
- Resources and skills: Appointment of a person for each of the following roles: Pilot leader, Platform administrator, Responsible for communication, Responsible for training and Responsible for collecting feedback.
- Communication to stakeholders: Definition of the way in which pilot partners will distribute the information to stakeholders, and how the stakeholders are involved.
- Communication to young citizens: Definition of the strategy and the means to be used in order to engage young citizens.
- Social media management: Definition of the social media to be used and description of usage rules.
- Inclusion: Plans to promote and ensure inclusion of sensitive communities.

The STEP platform was launched to the general public on February 2017 and since then STEP pilots have uploaded a number of questions on local environmental issues. The evaluation of the pilots will take place within the next months and will be concluded on October 2017.

4 The Technology Behind STEP

The STEP platform is a cloud based SaaS platform that brings together several ideas and technologies with the goal to engage Young European Citizens and Public Authorities in decision making about environmental strategies, policies, plans, programmes, laws, and projects.

The STEP platform is based on a modular service oriented architecture, integrating individual components which are developed/customized so that they can carry out specific business functions, and can be reused, each one individually, by public organizations for quickly opening their decision-making processes. This is especially important for public organizations that already have well set-up procedures for managing participation, and do not want to replace them. Public organizations can however benefit from the use of partial components of STEP, according to their needs, thus integrating them in their regular practice. STEP Platform has the following integrated components: e-Participation, Social media/Web mining and Visualization, Machine Translation and Text-to-Speech. STEP is available on the web at step.green, on Google Play and on App Store in English, Spanish, Italian, Greek, Turkish and Catalan languages.

4.1 e-Participation

The cloud e-Participation component is the central module of the STEP platform, allowing the interaction between end users (policy makers and young people) and the communication with the other platform components. e-Participation supports three types of dialogue management (e-Petition, Consultation, Idea Generation), Mapping, Timeline, Questionnaire, Messaging, Round Table, Embeddable Snippets and leaderboard and badged based Gamification features for public administrations and citizens. The STEP Platform follows the guidelines of Service Oriented Architecture (SOA) for the integration of all platform components and integration with 3rd parties including iOS and Android mobile apps via REST based services. e-Participation component uses Microsoft. Net technologies and applies web and mobile security practices to ensure the confidentiality, integrity and availability of electronic information captured, stored, maintained, and used by the STEP platform. It also provides a common logging service is being developed for as a user interaction log system within the platform.

4.2 Social Media/Web Mining and Visualization

STEP Platform uses social media and web mining technologies that enable young people to view personalized and enriched information relevant to the topic of their interest. Information is filtered through multiple social media and web sources, such as Facebook, twitter, and so on, and presented in a single unified environment. This tool eliminates the need to search for content on multiple platforms, or navigate through the overflow of information available. Web and social media mining component presents emerging topics from social media and the web relevant to users' interests in two different interfaces separately for PAs and citizens. For PAs, visualization of the content is provided via map and timeline based visualisations to help the analysts to identify the spatial distribution of use contributed content, the origin of certain trends and their evolution over time. Social media/Web mining and Visualization uses technologies such as Java, MongoDB, Solr, HTML5, jQuery and REST services for stream management, crawling and API management. Social media mining technologies are further elaborated in Sect. 5 of this paper.

4.3 Machine Translation (MT) and Text-to-Speech (TTS)

STEP Platform uses machine translation technology to allow users to access content from other countries in their own language (English, Spanish, Italian, Greek, Turkish and Catalan), removing language barriers and fostering cross-border interaction.

Statistical Machine Translation (SMT) is used in STEP MT that is based on on automatic corpus analysis and statistical learning methods, where both translations and the rules are learned from a sufficient amount of bilingual corpora, i.e. from existing translations. Besides MT, Text to Speech (TTS) technology is used with broad functionalities to allow flexible access from the STEP platform. TTS enables dynamic content provided by users to be read in English, Spanish, Italian, Greek, Turkish and Catalan languages. TTS can convert any text into an WAV and MP3 audio file formats with the following voice parameters support: Speed, Pitch, Volume and Gender.

5 Social Media Mining

Social media mining tool is a generic open-source framework for monitoring, analyzing and retrieving content form multiple social media platforms.

Figure 1 depicts the components of the tool, deployed as separate Docker containers. End users of the tool can define collections that represent their information needs through a REST API or by using a usr interface on top of it. These collections can contain (1) keywords, (2) accounts or (3) locations. The tool uses that input to collect relevant content and expose analytics over it.

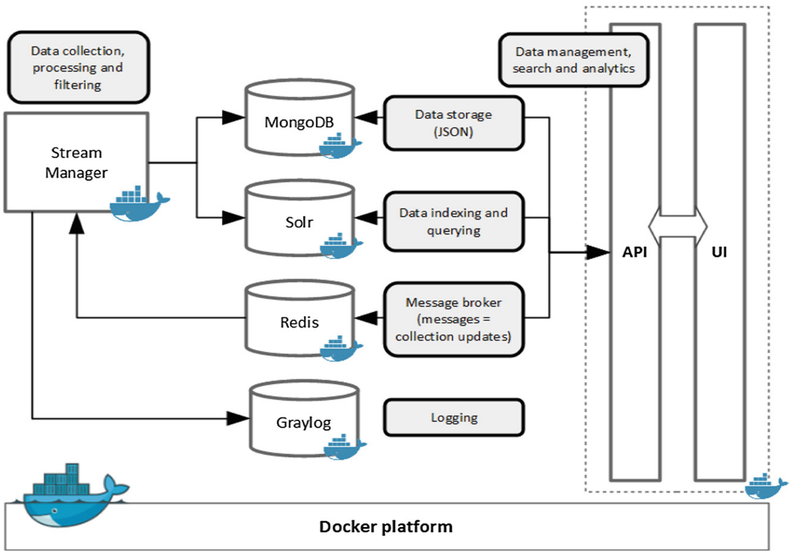

Fig. 1. Overview of social media mining tool and its core components

The collection of social media data is performed by the Stream Manager component. Signaling of Stream Manager for collection related events occurred through the API is performed by the Redis publish/subscribe mechanism. Currently, five different platforms are supported: Twitter, Facebook, YouTube, Google+ and Flickr. The Stream

Manager collects items shared in these platforms, alongside the users that published them, the embedded multimedia and the linked URLs. The fetched content, after being processed and filtered to improve quality, is stored and indexed in MongoDB and Apache Solr respectively.

Given a user defined collection retrieval of relevant content is performed by using a Solr query. Solr response contains the ids of relevant items, while the actual content is retrieved from MongoDb by id. For analytics, faceting mechanism of Solr is used. More precisely, using the same Solr query as content retrieval facets are calculated over several fields such as publication time, user ids, tags etc. These facets correspond to the analytics provided by the tool.

6 Conclusions

This paper presents STEP project, a H2020 project with basic aim to engage even more young people in societal and political participation and to involve them in decision making processes on issues that concern the environment. In this context an interactive platform for e-participation has been developed, which addresses mainly two target groups: young citizens and the representatives of the authorities, who are responsible for decision making. Using the latest advances of technology, the STEP platform has been enriched with a number of features that make the platform friendlier to the user. The STEP platform is currently being pilot tested by 6 authorities in four countries.

References

1. Public Opinion in the European Union, Standard Eurobarometer 77. Spring 2012, p. 13 (2012)
2. Potential and Challenges of E-Participation in the European Union, study commissioned and supervised by the European Parliament's Policy Department for Citizens' Rights and Consti-tutional Affairs at the request of the AFCO Committee (2016)
3. e-Participation Best Practice Manual, This manual is the result of Interreg IVC-financed project "eCitizen II – Towards citizen-centered eGovernment in European cities and regions (2012)
4. EUth project. http://www.euth.net/. Accessed July 2017
5. Environmental Protection Agency, Public Participation Guide. https://www.epa.gov/sites/production/files/2014-05/documents/ppg_english_full-2.pdf. Accessed 24 July 2017
6. Victorian Auditor-General's Office. https://www.audit.vic.gov.au/sites/default/files/20150130-Public-Participation-BPG.pdf. Accessed 24 July 2017
7. International Youth Service of the Federal Republic of Germany Guidelines for successful e-participation by young people in decision-making at local, regional, national and European levels
8. International Association for Public Participation. http://www.iap2.org/?page=a4. Accessed 24 July 2017
9. National Coalition for Dialogue and Deliberation. http://ncdd.org/rc/wp-content/uploads/2010/08/PEPfinal-expanded.pdf. Accessed 25 July 2017

10. World Bank. http://web.worldbank.org/WBSITE/EXTERNAL/TOPICS/EXTSOCIALDEVELOPMENT/EXTPCENG/0,,contentMDK:20282087~menuPK:1278110~pagePK:148956~piPK:216618~theSitePK:410306,00.html. Accessed 24 July 2017
11. Environmental Protection Agency. https://www.epa.gov/international-cooperation/better-decisions-through-consultation-and-collaboration. Accessed 25 July 2017
12. International Association for Public Participation. http://c.ymcdn.com/sites/www.iap2.org/resource/resmgr/foundations_course/IAP2_P2_Spectrum_FINAL.pdf. Accessed 24 July 2017
13. Environmental Protection Agency. https://www.epa.gov/international-cooperation/how-consult-and-involve-public. Accessed 24 July 2017
14. Environmental Protection Agency. https://www.epa.gov/sites/production/files/2015-09/documents/evaluate508.pdf. Accessed 24 July 2017

Understanding and Empowering Digital Citizens

A Study of Ride Sharing Opportunities in the City of Santiago de Chile

Emanuele Massaro[1,2](✉)

[1] Laboratory for Human-Environment Relations in Urban Systems (HERUS), Swiss Mobiliar Chair in Urban Ecology and Sustainable Living, Institute of Environmental Engineering (ENAC), École Polytechinque Fédérale de Lausanne (EPFL), 1015 Lausanne, Switzerland
emanuele.massaro@epfl.ch

[2] Dipartimento di Ingegneria dell'Informazione, Università Degli Studi di Firenze, Via Santa Marta, 3, 50139 Firenze, Italy

Abstract. As an industrial revolution booms in Chile, the country's air has been flooded by toxic emissions. Urban cities face the worst of the pollution, as factories are booming and urban centers are growing. Indeed, one of the main contributors towards the accumulation of PM2.5 are cars. Sharing trips may help reduce the number of private and public vehicles on the road and cut down on greenhouse gas emissions, travel time and individual costs. In this research I apply the concept of shareability networks to a survey of 113,591 of trips taken in the city of Santiago in Chile, showing that with increasing but still relatively low passenger discomfort, cumulative trip length can be cut by 50% or more. I quantify the benefit of ride sharing in terms of traffic and emission reduction. I finally show that the ride sharing potential is substantial, with nearly 100% of the trips shareable with current public transportation trip demand.

Keywords: Shareablity network · Complex systems · Santiago de Chile

1 Introduction

Santiago is one of the most prosperous and advanced Latin America capitals, but also one of the most polluted cities in the Region. The Santiago metropolitan area experienced a fast growth in the last decades, from about 4 million of inhabitants in 1992 to the current 6.7 millions: Fast growth and urban sprawl have caused increasing environmental concerns and among these, air pollution represents a severe issue. The geographic position of the city, lying in a valley surrounded by high mountains, makes the air pollution problems acute, with periods of very high concentration of PM 10 and PM 2.5. Thus in the last years many measures have been implemented in Santiago to better understand and subsequently reduce air pollution in the city, and there is clear evidence that transportation represents one of its main causes. The goal of this study is providing a quantitative estimation of the potential benefits of ride sharing in the extended Santiago metropolitan area (Fig. 1). The study will also assess the potential environmental benefit due to shared mobility.

I. Kompatsiaris et al. (Eds.): INSCI 2017, LNCS 10673, pp. 159–173, 2017.
https://doi.org/10.1007/978-3-319-70284-1_13

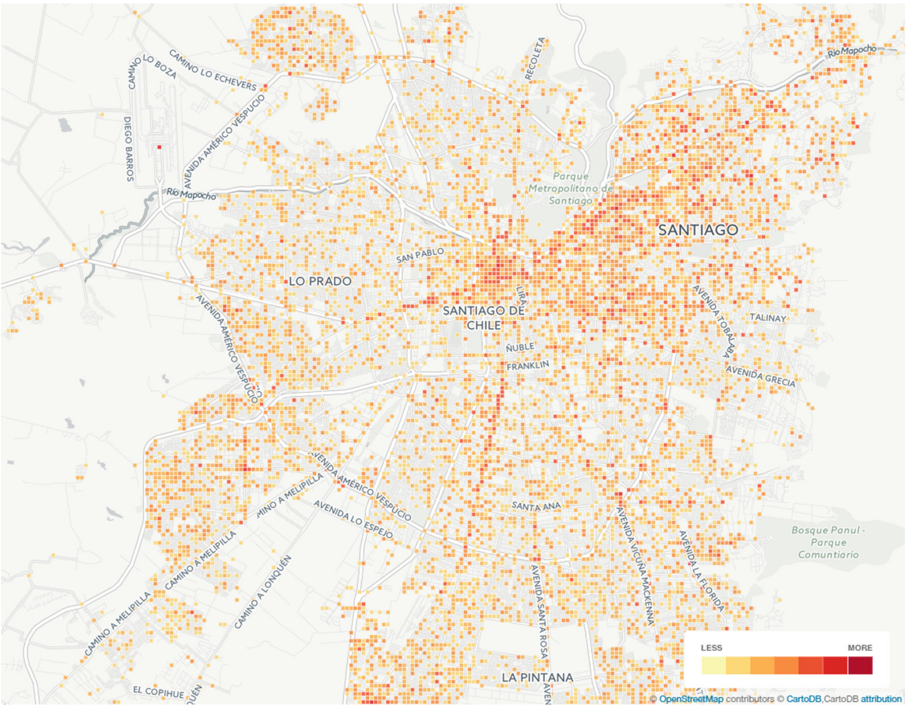

Fig. 1. Points of origin of the communing travels in the city of Santiago extracted from the SECTRA database.

2 Data Description

The study is based on data extracted from an extensive mobility survey conducted in 2012, and made publicly available by SECTRA (Secreteria del Planificacion de Transporte – Ministerio de Transporte y Telecomunicacion) through a website[1]. The survey has been conducted over an area of the Santiago Metropolitan Region (or simply Metropolitan Region) one of Chile's 15 first-order administrative divisions. The surveyed area includes the full Santiago Metropolitan Area, (32 comunas, i.e. municipality) and the neighboring areas populated by a relevant number of commuters, for a total number of 45 comunas. The area hosts a population of approximately 6.8 million people, representing more than 90% of the population of the Santiago Metropolitan Region, with an estimated number of 1.160.000 private vehicles, 6300 urban buses, 27.000 taxi, 11000 taxi colectivos (sharing taxi) and 5 metro lines with more than 104 Km of lines. Main aim of the survey was to collect detailed information on the trips daily performed in the selected area considering an appropriately selected sample of 18000 households. The collected information includes timing of the trips, path and transportations modes (public/private). The mobility survey comprises records of 113,591 trips

[1] http://www.sectra.gob.cl/index.htm.

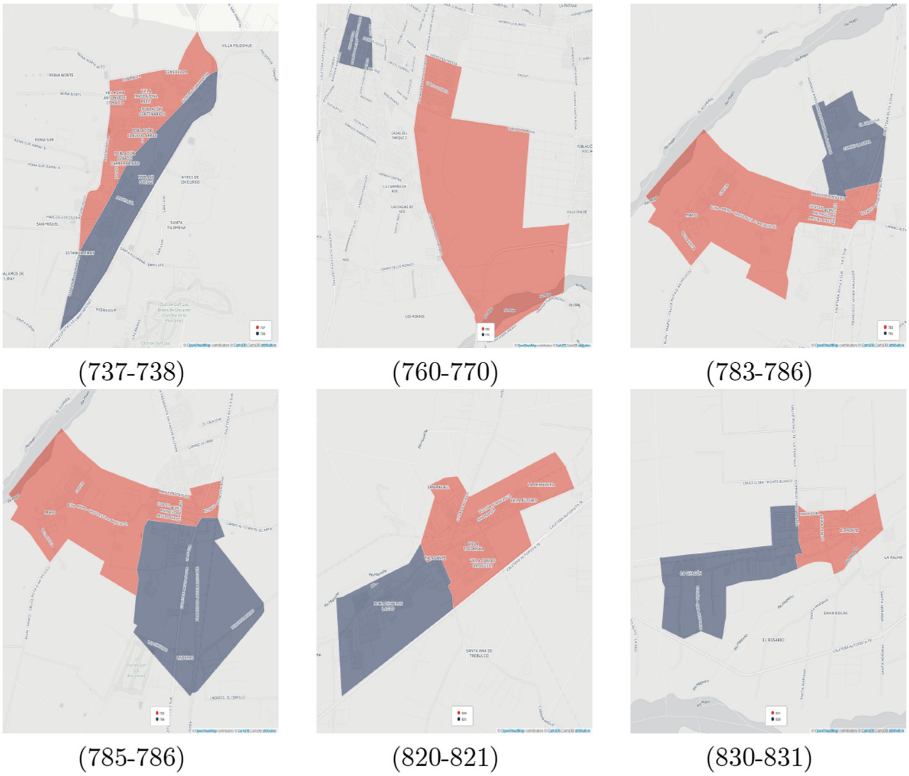

Fig. 2. Maps of the most important zones in the terms of total number of trips. I show the top 12 zones for what concerning external trips.

performed in the above described extended Santiago metropolitan area. The area is divided in zones, corresponding to a zonification process referred to different urban and socioeconomic parameters. Figure 2 reports the general view of the zones. Each trip record reports, among other things, information about the origin and destination of each trip, characterized with the granularity of the zones, the starting time of the trip, and transport mode. This latter information is used to select the subset of trips relevant for our study, namely, trips performed using either a private vehicle or a taxi. These amount to about 31,376 trips, which are still a representative sample of mobility in Santiago. In particular I consider only the categories 1, 4, 5, 6, 8, 11, 12 of the Modo Agregado field, as reported in Table 1.

Unless otherwise stated, this is the reference data set considered in the remainder of this study. These 31,376 trips are relatively shorts: if I observe the linear distance between two origin destination points in fact the 95% of them are below 18 km with an average length of 5.75 km. In Fig. 7(a)–(b) I report the distribution of the trip distance while in Fig. 8(c)–(d) I report the distribution of the trip duration. The average trip duration is 1712 s, and 95% of the

Table 1. Different categories of means of transportation in our database. The dataset is composed by 18 sub-categories for the categories *Modo* and *Modo Agregado*, 6 for the category *Modo Difusion* and 3 sub categories for *Modo Motor*. Here *Modo* means *Mean of transportation*, *Modo Agregado* means *Aggregated mean of transportation*, *Modo Difusion* means *Way of moving* and *Modo Motor* tells if the way of moving has a motor or not.

Category	Modo	Modo Agregado	Modo Difusion	Modo Motor
1	Private Driver	Car	Car	Motor Yes
2	Urban Bus	Bus TS	Bib!	Motor No
3	General Bus	Bus no TS	Sharing Taxi	Other
4	Metro	Metro TS	Taxi	
5	Sharing Taxi	Sharing Taxi TS	Bib! - Other Public	
6	School Bus as passenger	Taxi	Bib! - Other Private	
7	Taxi or Radiotaxi	Bus TS - Bus no TS	Other	
8	Walk	Auto - Metro	Walk	
9	Bicycle	Bus TS - Metro	Bicycle	
10	Motorbike	Bus no TS - Metro		
11	Institutional Bus	Sharing Taxi - Metro		
12	Bus Rural	Taxi- Metro		
13	School Bus as Escort	Other - Metro		
14	Metrobus	Other - Bus TS		
15	Informal Services	Other - Bus TS - Metro		
16	Train	Sharing Taxi - Other		
17	Auto Escort	Taxi- Walk		
18	Motorbike Escort	Bicycle		

trips are below 80 min. When interpreted as the probability $P(r)$ of finding a displacement of length r in a short time dt, the existence of a typical length scale often justifies the description of dispersal in terms of diffusion equations on large spatio-temporal scales. If, however, $P(r)$ lacks a typical length scale, that is $P(r) \sim r^{-(1+\beta)}$ with $\beta < 2$ as in this case, the diffusion approximation fails. In physics, random processes with such a single-step distribution are known as Lévy flight [1,2].

In order to evaluate the real distance of a travel I made use of the Google Maps APIs[2] which allows us to extract the real distance between two points A and B given the coordinates expressed in Latitude/Longitude UTM format. In this case I observed an average communing travel distance equal to 16.32 km while 95% of the trips are below 40 km with median 12 km as reported in Fig. 8 in the Appendix. From this point on, I will make use of the distance extracted from the Google Maps API for our analysis. I observed that most of the trips appear in the same zone (see Table 3 in the appendix): in Fig. 7 I report the map of the top 10 zones that count the higher number of internal trips – i.e., trips with origin and destination in the same zone that are always in some comuna as reported in Table 2. In a similar way in Fig. 2 I show the most important

[2] https://developers.google.com/maps/.

Table 2. Correspondece between Zona (Zone) and Comuna (Municipality).

Zona (Zone)	737	738	760	770	783	786	785	820	821	830	831
Comuna (Municipality)	76	76	98	98	103	103	103	86	86	89	89

connecting zones in terms of extra-zones trips. Here I take into account only the commuting trips between two different zones (Fig. 11). The table of the top 11 extra communing zones is reported in the appendix (Table 4), while the general map is reported in Fig. 10.

3 Shareability Network

Sharing opportunities between any two trips are determined on zone-based approach that exploits subdivision of the Santiago Metropolitan area into non-overlapping zones. The complete map of the origin destination points is reported in the Appendix in Fig. 9. Each trip record reports also the origin and destination zone of a trip. Two trips are considered to be shareable if they have the same origin and destination zone, and if their starting times are at most T minutes apart, where T is a tunable parameter used to express user tolerance to ride sharing. Sharing opportunities between any possible trip pair will be determined based on the approach described above, with the goal of building the corresponding shareability network (Fig. 3) as defined in by Santi et al. [6]. The shareability network is a mathematical model of ride sharing opportunities that allows efficient computation of the optimal ride sharing solution.

Once the shareability network is formed, ride sharing opportunities will be quantified according to the following metrics:

- Percentage of shareable trips, i.e., percentage of trips that have at least one sharing opportunity. This corresponds to the fraction of nodes in the shareability network that have degree at least one.
- Average number of sharing opportunities, i.e., the average number of sharing options available for each trip. This corresponds to the average node degree in the shareability network.
- Shareability, i.e., fraction of trips that can be shared. This metric is defined by first computing the maximum matching on the shareability network, which corresponds to the maximum number of pair-wise trip sharing possible, then counting the number M of nodes that belong to the maximum matching, and finally dividing M by the total number N of nodes in the shareability network.

I count 17860 distinct locations A and B where at least one trip was made. However, in order to build the shareability network I need to take into account only the possible combinations with more than one trip: 5167 possible combinations contain at least 2 trips. Then in order to take into account the shareable effect I consider the all the possible combinations to share a trip between

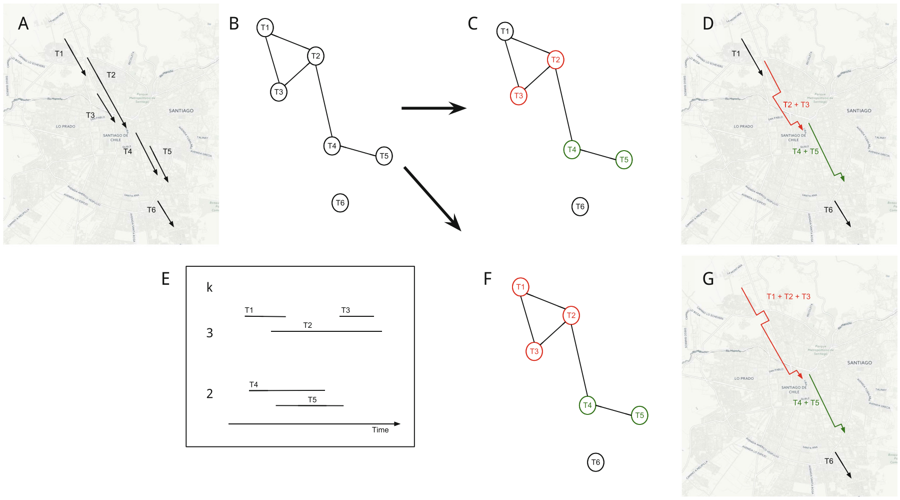

Fig. 3. Shareability networks translate spatiotemporal sharing problems into a graph-theoretic framework that provides efficient solutions. (A) Example of six trips, $T1, \ldots, T7$, requested and to be shared in Santiago. (B) Construction of shareability network for $k = 2$. Trips that could potentially be shared are connected, given the necessary time constraints to hold which I assume here to be the case. Trip 6 is an isolated node because it cannot possibly be shared with other trips. (C) Maximum matching of the shareability network gives the maximum number of trip pairs, i.e., the maximum number of shared trips. (D) Implementation (routing) of the maximum matching solution. (E) Each of the two cases involves a number of trips T_i to be shared, but ordered differently in time t. (E) Alternative implementation of solution (B) with different shareability time. (F) Implementation (routing) of the alternative solution.

a zone A to a zone B. In particular let say that if we have 4 shareable trips between A and B, so $t^r_{a,b} = 1, 2, 3, 4$, then all the possible combinations are $c_{a,b} = [(1, 2), (1, 3), (1, 4), (2, 3), (2, 4), (3, 4)]$. In general the total number of combinations $t^c_{a,b}$ between two zones A and B is given by:

$$t^c_{a,b} = \frac{1}{n_{a,b}} n_{a,b}(n_{a,b} - 1), \quad (1)$$

where $n_{a,b}$ is the total number of shareable trips between A and B.

I counted 58611 possible combinations but only 11596 if we consider a shareability time $T = 30$ min. In order to detect the most interesting shareability zones, I consider zones that share more than 10 trips with a sharing frequency of at least the 50%. The sharing frequency between two zones A and B is computed as the fraction between the shareable and the total trips. As result of this assumption I detect 20 zones as reported in Fig. 4 where I show the comparison between shareable trips and sharing frequency of the aforementioned zones as reported in Fig. 5.

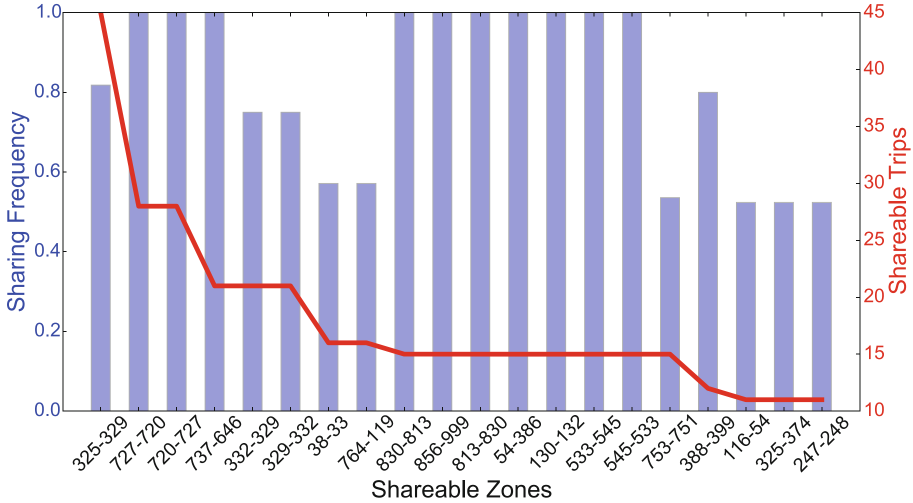

Fig. 4. Comparison between sharing frequency (blue) and shareable trips (red) in the zones with more than 10 shareable trips and at the least the 50% of sharing frequency. (Color figure online)

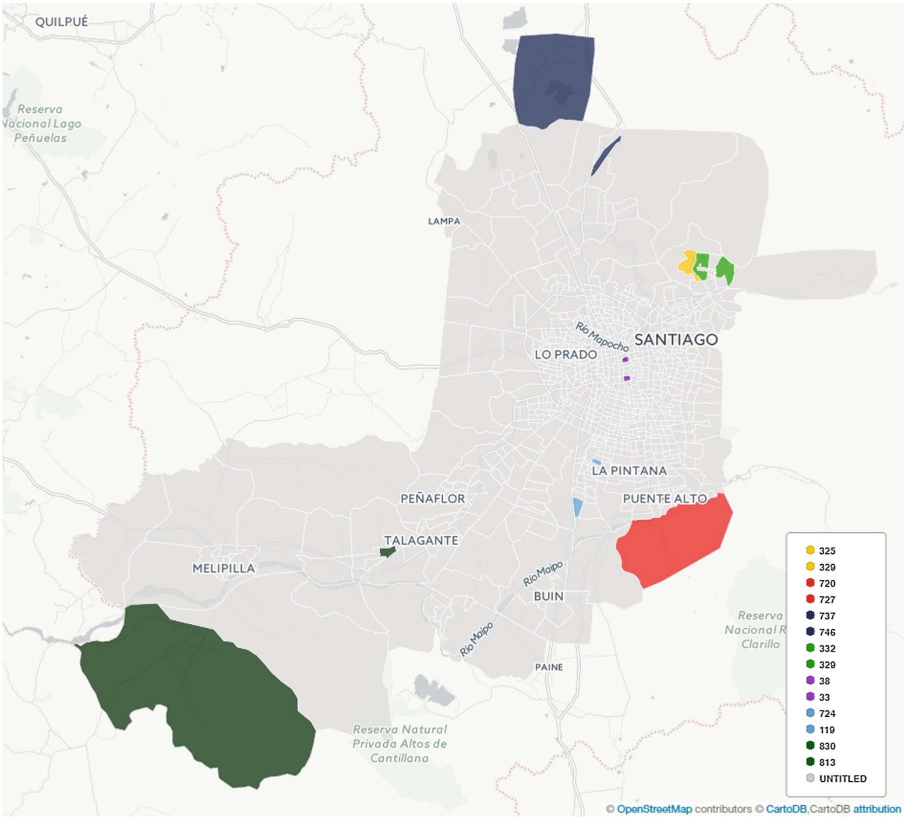

Fig. 5. Maps of some of the shareable zones reported in Fig. 4.

4 Quantifying Emission Reduction

Here I show the effect of the shareability in terms of emissions reduction. In particular I estimate the total amount of kilometers made from all the trips in our database. In doing so I count all the possible trips among the 5167 coupled ones. In total I count 19139 trips and 143749 km. If we consider a shareability time of 30 min we count 2510 shareable trips that allow us to save more than 20587 Km that represent the 14% of the distance among the possible shareable coupled zones (which is around $29 \cdot 10^4$ km). To translate the above traveled distance reduction into emission reduction, I consider the following well-to-wheels emission factors for gasoline vehicles [3]: 0.05 g/km for PM_{10}, 0.024 g/km for $PM_{2.5}$, and 2.25 g/km for carbon monoxide (CO_2). Total daily emission reductions then amount to 1.029 Kg of PM_{10}, 0.494 Kg of $PM_{2.5}$, and 46.32 Kg of CO_2, summing up to 375.58 kg, 180 Kg, and 16.906 tonnes, respectively, for the whole year. Even more substantive emission reductions could be achieved if shared trips would be served by a fleet of electric shared vehicles. In this case, the total traveled kilometers of the 2510 shared trips, amounting to km, would imply zero emission, and a yearly emission reduction of 14.5 tonnes of PM_{10}, 6.96 tonnes of $PM_{2.5}$, and 652 tonnes of CO_2, respectively. Please notice that these emissions savings, already substantial, account only for the trips comprised in the mobility survey, which represent a small fraction of the actual trips in the city of Santiago. We can then expect much larger savings in case ride sharing coupled with electric mobility would be extended to the entire mobility demand in Santiago. However, due to lack of suitable data sets, at the current status it is not possible to quantify these reductions with a scientifically sound methodology.

5 Conclusions

In this study I analyzed the effect of public ride sharing in the city of Santiago. I quantified the reduction in terms of traveling and in terms of emission. The mathematical tool I used to draw optimal sharing routes is called “shareability networks” [6]. The key idea was to express the shareable resource, in this case, a taxi trip, as a node in a network, and to draw links between two nodes if the corresponding resources – trips – could be shared. A recent study has shown that taxi ride shareability – i.e., the fraction of taxi trips that can be shared – can be well approximated by knowing a number of city parameters [7]. Namely, city area, average speed of traffic, and taxi demand. With a city area of approximately 641 km^2 and an average traffic speed of 31 km/h, the taxi ride sharing potential in Santiago is reported in Fig. 6. As seen from the figure, the ride sharing potential is substantial, with nearly 100% of the trips shareable with current taxi trip demand. This benefit comes with reductions in service cost, emissions, and with split fares, hinting toward a wide passenger acceptance of such a shared service. Such reductions in car numbers would dramatically lower

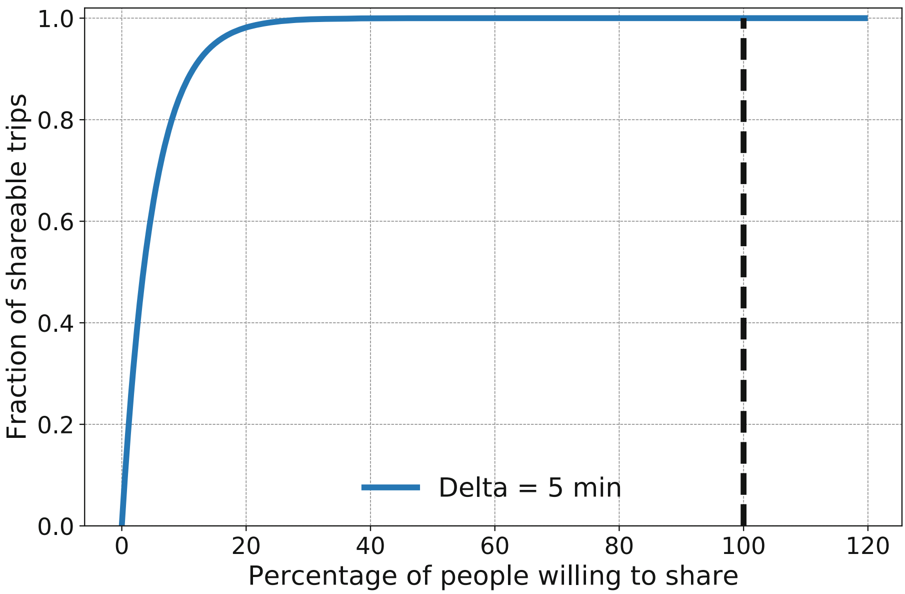

Fig. 6. Taxi ride sharing potential in the city of Santiago, estimated according to [7], assuming that shared rides incur a delay Δ of at most 5 min. With current taxi demand (dashed line), nearly 100% of the trips can be shared. Ride sharing potential remains high even if one in five customers opts for sharing a taxi ride.

the cost of our mobility infrastructure and the embodied energy associated with building and maintaining it. Fewer cars may also mean shorter travel times, less congestion, and a smaller environmental impact [5]. Future studies should see the development of a model that includes more social and economic factors in shareability, such as safety (whether riding with a stranger might be perceived as dangerous) and pricing. Moreover we should take in consideration about possible unintended consequences, such as lost jobs for taxi drivers or cuts in bus services [4]. Already, some US cities are reportedly considering cuts to the availability of public transportation, under the assumption that stranded customers can instead use a ride-hailing company.

Acknowledgments. I would like to thank Renata Mele and Daniela De Rosa from Enel Foundation for their support and assistance with this research study. I would also like to thank all the members of the MIT Senseable City Lab consortium for partially supporting this research and for providing the data for generating Fig. 6. I would like to thank the DISIT Lab of the Department of Information Engineering of the University of Florence for partially supporting this research.

Appendix

Table 3. Top 10 intra-zones in terms of total number trips.

Origin/destination zone	Comuna (municipality)	Trips
727/727	Pirque	77
738/738	Colina	70
325/325	Lo Barnechea	60
786/786	Buin	56
821/821	Talagante	46
398/398	Maipù	46
150/150	Huechuraba	44
307/307		32
820/820		20
490/490		27

Table 4. Top 11 extra-zones in terms of total number trips.

Origin/Destination zone	Comuna (municipality)	Trips
738/737	Colina	56
783/786	Buin	42
760/770	San Bernardo	42
820/821	Talagante	38
770/760	San Bernardo	38
785/786	Buin	38
786/785	Buin	37
786/783	Buin	35
831/830	El Monte	34
821/820	Talagante	34
830/831	El Monte	32

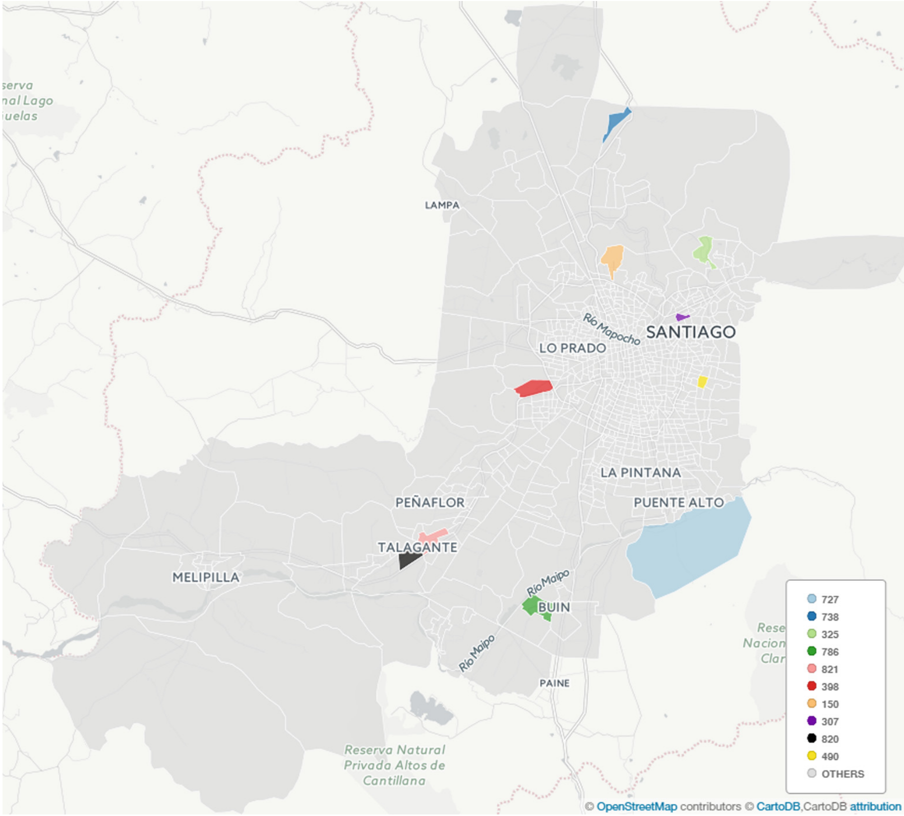

Fig. 7. Maps of the most important zones in the terms of total number of trips. I show the top 10 zones for what concerns internal trips.

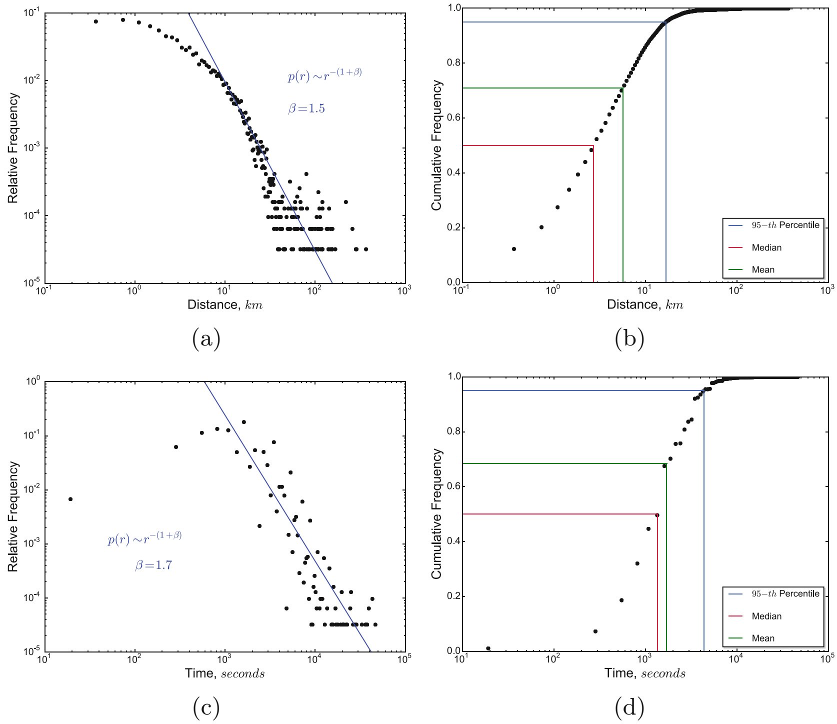

Fig. 8. Frequency distribution of the communing time and the linear distance between origin destination points.

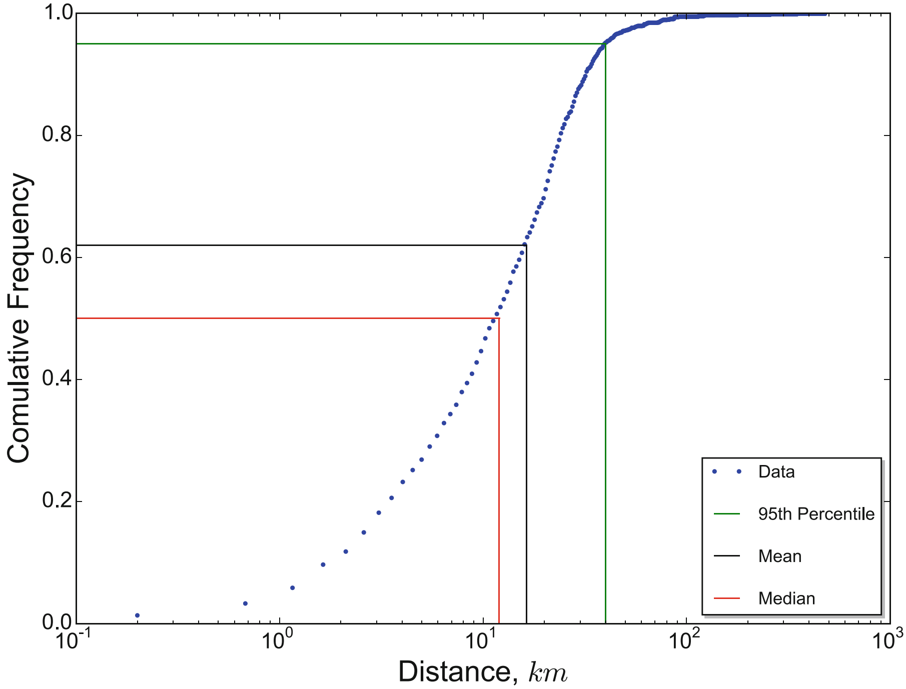

Fig. 9. Cumulative frequency distribution of the travel distance extracted from Google Maps APIs between origin destination points.

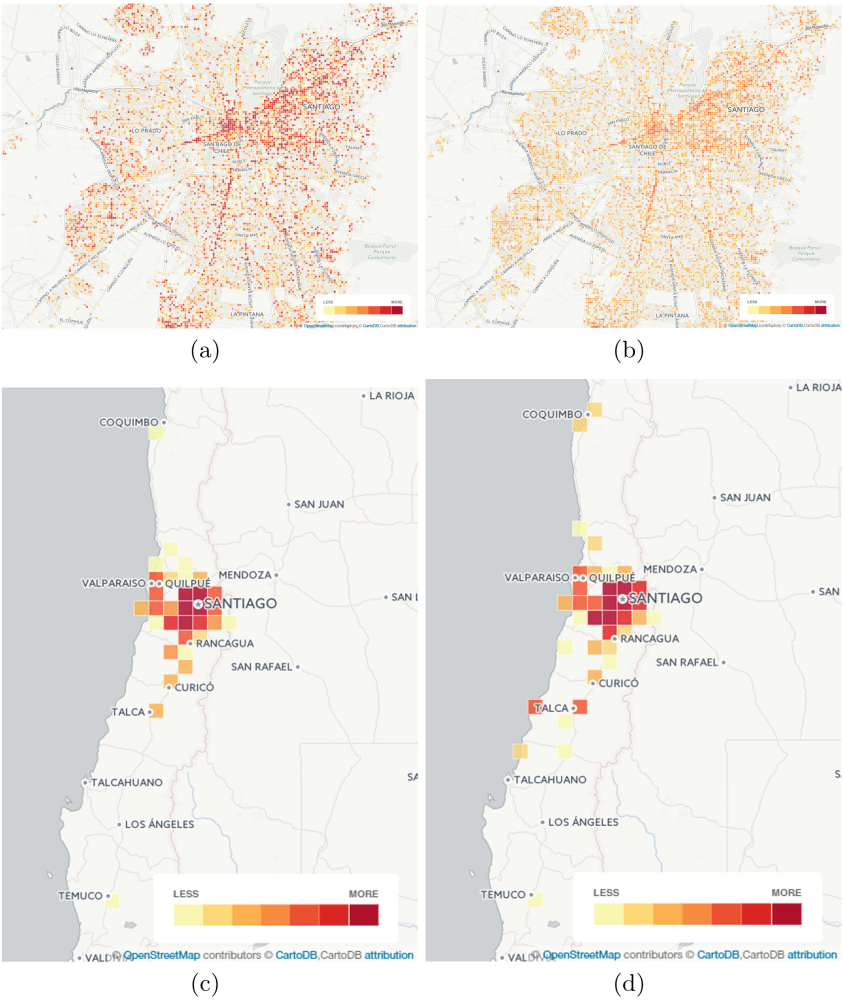

Fig. 10. Origin destination points in the area of city of Santiago (a)–(b) and in the gran Metropolitana Aerea (c)–(d).

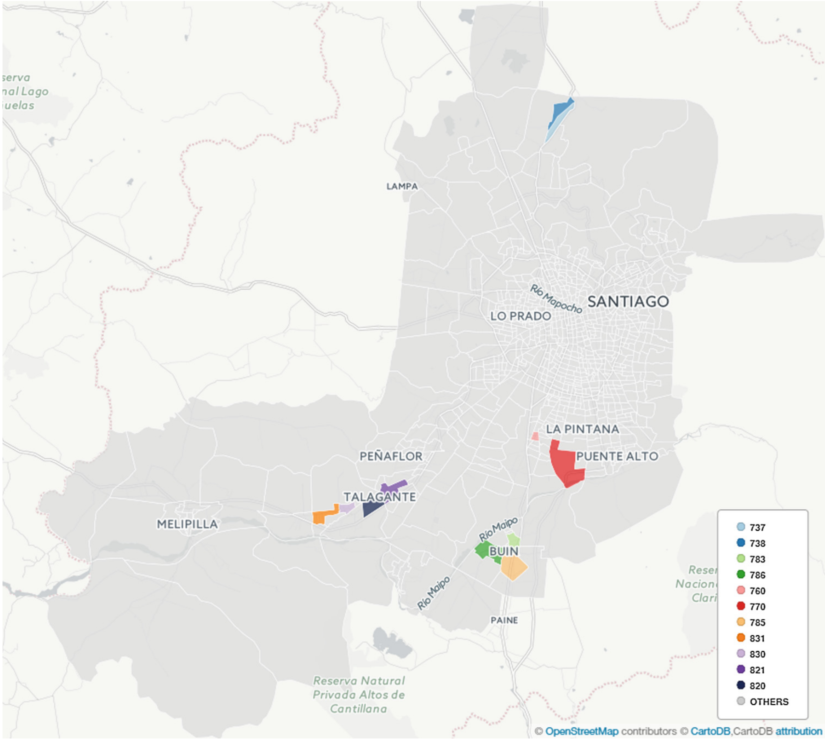

Fig. 11. Most important zones in terms of total extra-zones communing trips.

References

1. Brockmann, D., Hufnagel, L., Geisel, T.: The scaling laws of human travel. Nature **439**(7075), 462–465 (2006)
2. Gonzalez, M.C., Hidalgo, C.A., Barabasi, A.L.: Understanding individual human mobility patterns. Nature **453**(7196), 779–782 (2008)
3. Huo, H., Wu, Y., Wang, M.: Total versus urban: well-to-wheels assessment of criteria pollutant emissions from various vehicle/fuel systems. Atmos. Environ. **43**(10), 1796–1804 (2009)
4. Martin, C.J.: The sharing economy: a pathway to sustainability or a nightmarish form of neoliberal capitalism? Ecol. Econ. **121**, 149–159 (2016)
5. Massaro, E., Ahn, C., Ratti, C., Santi, P., Stahlmann, R., Lamprecht, A., Roehder, M., Huber, M.: The car as an ambient sensing platform [point of view]. Proc. IEEE **105**(1), 3–7 (2017)
6. Santi, P., Resta, G., Szell, M., Sobolevsky, S., Strogatz, S.H., Ratti, C.: Quantifying the benefits of vehicle pooling with shareability networks. Proc. Nat. Acad. Sci. **111**(37), 13290–13294 (2014)
7. Tachet, R., Sagarra, O., Santi, P., Resta, G., Szell, M., Strogatz, S.H., Ratti, C.: Scaling law of urban ride sharing. Sci. Rep. **7**, 42868 (2017). http://dx.doi.org/10.1038/srep42868

Smart Cities in Stars: Food Perceptions and Beyond

Maria Pontiki(✉), Panagiota Koltsida, Dimitris Gkoumas, Dimitris Pappas, Haris Papageorgiou, Eleni Toli, and Yannis Ioannidis

Athena Research and Innovation Center in Information, Communication and Knowledge Technologies, Artemidos 6 & Epidavrou, 15125 Maroussi, Athens, Greece
{mpontiki,dgkoumas,dpappas, xaris,yannis}@athena-innovation.gr, {p.koltsida,elto}@di.uoa.gr

Abstract. Citizens are shaping their food preferences and expressing their food experiences on a daily basis reflecting their way of living, culture and well-being. In this paper, we focus on food perceptions and experiences in the context of smart citizen and tourist sensing. We analyze Foursquare user reviews about food-related points of interest in ten European cities, and we explore the imprint of a city as it is shaped based on the spatial distribution of food-related topics and sentiments. The topic modelling and sentiment analysis results are visualized using geo-referenced heat maps that enrich the cities maps with information that allows for a more insightful navigation across their different geographical regions providing insights not available in the original data.

Keywords: Urban sensing · Food experiences · Topic modelling · Sentiment analysis

1 Introduction

The emergence of social networks, content publishing platforms, check-in applications and smartphones/GPS devices in recent years has created huge amounts of user-related data. To exploit these rapidly growing data, recent research has focused on the use of the geographic side of information to trace events [1], analyze the sentiment of users [2], identify popular Points of Interest (POIs) [3], identify and visualize typical day-to-day traffic patterns [4], as well as to improve existing city maps. In this context, the growing need for cost-effective location-based services and effective online advertising has led leading online service providers (such as Google, Bing, Foursquare) to store and distribute information about POIs to users of these services, usually through the use of web-based programming interfaces. Users not only have access to information about nearby POIs, but they also have the ability to read and provide reviews or to inform their friends about their current location.

Foursquare data have been explored in the context of location based social networks and urban sensing in different ways to explore research challenges such as

I. Kompatsiaris et al. (Eds.): INSCI 2017, LNCS 10673, pp. 174–183, 2017.
https://doi.org/10.1007/978-3-319-70284-1_14

identifying the different factors which drive check-in behavior by clustering users into meaningful groups using topic modelling techniques [5], investigating the impact of tips on user mobility using sentiment analysis techniques [6], or building location recommendation models by combining the preference extracted from check-ins and text-based tips which are processed using sentiment analysis [7].

The work presented in this paper aims at contributing to smart citizen and tourist sensing with a focus on food perceptions and further insights that can be derived from user reviews. To this end, we have been collecting and analyzing user reviews about food-related POIs in ten European cities, and we explore the imprint of each city as it is shaped based on the spatial distribution of food-related topics and sentiments discussed by food consumers. The results are visualized using geo-referenced heat maps enriching the cities maps enabling answers to questions like *"Burger food options in downtown Athens*" or "*Good Breakfast options close to Louvre Museum in Paris*". The comparative analysis of the results for the ten cities can provide insights like "*In which European city the service is mostly evaluated as friendly*?".

Moreover, capturing and analyzing information about people's food choices and eating behavior can help consumers, food providers, and policy makers in making informed decisions on their activities (e.g. food recommenders, targeted nutrition education programming). The work presented in this paper constitutes a part of an ongoing work towards this direction and is developed through the CAPSELLA Social Data Platform. The CAPSELLA[1] project aims to support communities of farmers and food manufacturers in making informed decisions on their activities by offering them access to data from a variety of open data sources related to regional agro-biodiversity and food. A core activity towards this goal is the development of Social Data services based on the needs and interests of a community, depending on their respective roles (i.e. consumer, producer, etc.). A brief overview of the CAPSELLA infrastructure and Social Analytics Platform is provided in Sect. 2. Section 3 presents the topic modelling and sentiment analysis methodology. The results are presented in Sect. 4. The paper concludes with a discussion on the results and further extensions of our work.

2 The CAPSELLA Social Analytics Platform

The CAPSELLA infrastructure (Fig. 1) is the base layer of the services and applications developed in the context of the CAPSELLA project and is designed to meet various requirements coming from different communities, user profiles and to support distinct needs. It consists of several independent, but highly inter-connected systems and offers a set of functionalities covering data and metadata management and data analytics, thus supporting the complete cycle of a data infrastructure. In addition, a number of interoperable services have been developed to make the infrastructure able to interact and exchange datasets and information by exploiting, among others, well known standards. The core systems of the infrastructure are: (a) the Data Management System, offering storage, retrieval and management facilities for various data types

[1] http://www.capsella.eu/.

including non-relational, tabular, relational and geospatial data; (b) the Metadata Catalogue System, which offers browsing and discovery of the datasets based on a set of descriptive metadata, and (c) the Data Analytics System (Social Data Platform) that lies next to the above systems and supports data analysis and processing, as well as decision making for its client applications. Furthermore, a central horizontal service has been developed so as to manage the authentication and authorization aspects of the infrastructure concerning both the users and the datasets. Extensibility, scalability and performance are the three pillars, on which the design and the implementation of the above systems are settled. In this direction, all the systems and services can be easily extended to meet additional requirements that were not considered in the initial design.

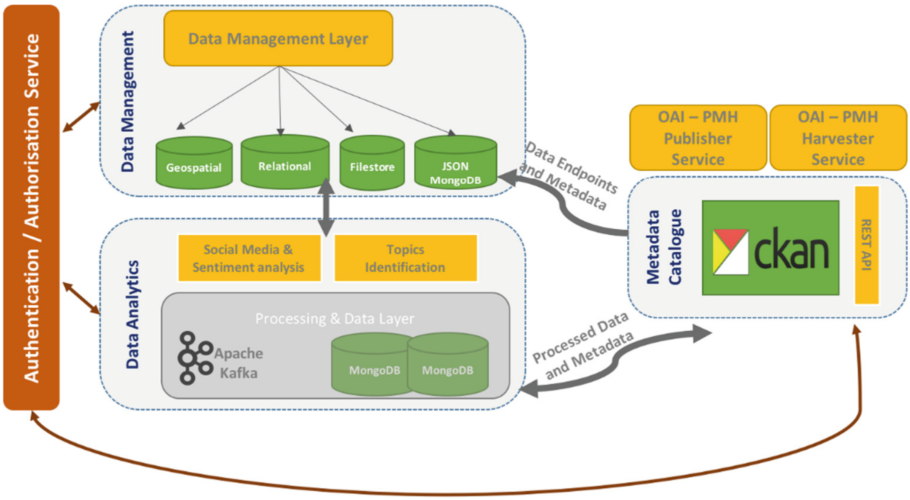

Fig. 1. The CAPSELLA infrastructure.

The CAPSELLA Social Data Platform consists of five major tiers. The Data collection tier where datasets are being captured or ingested from a variety of sources through a set of dedicated tools. Transporting data from the Collection tier to the CAPSELLA platform is facilitated by a message queuing tier (MQ tier) following the Publish/subscribe interaction pattern. In the CAPSELLA MQ tier, social analytics workflows (processing tier) play the role of the receivers enabling the efficient and effective processing of the data. In addition, data can be distributed within the platform to provide additional protections against failures, as well as significant opportunities for scaling performance. The results (e.g. annotations) are transmitted as new messages and further consumed by (a) a distributed indexing & visualization framework providing search and browsing functionalities (Data Exploration tier), (b) archiving to a storage system, and (c) a persistent storage subsystem providing REST APIs to web apps, third applications and the CAPSELLA pilot projects (access tier).

The social analytics workflows integrate several natural language processing and text analytics modules and pipelines that explore our daily food experiences, diversity in food behavior and culture, etc. and produce a wealth of annotations (e.g. sentiments, topics, entities like nutrients, food ingredients, restaurants etc.) enriching the original data collections. The workflows are designed based on the specifications and the requirements of each pilot in order to provide the right type of information to the right community of people. An example of such a workflow is the Stevia[2] Application that performs sentiment analysis on Twitter data focusing on specific stevia aspects (e.g. taste, nutrition value, healthiness and price). In this paper, we present the topic modelling and sentiment analysis workflows developed in the context of smart citizen and tourist sensing by focusing on food perceptions and further insights that can be derived from user reviews.

3 Methodology

In order to explore the imprint of a city as it is shaped based on experiences expressed by food consumers, we collect and analyze user reviews about food-related POIs in ten European cities. The research hypothesis is that the spatial distribution of food-related topics and sentiments discussed by food consumers can provide to citizens and tourists meaningful information not available in the original data. In this context, we employ topic modelling, an unsupervised machine learning method, to detect topics of interest. In a second phase, we perform sentiment analysis to capture sentiments expressed towards the extracted topics for each city.

3.1 Data Collection

We collected from Foursquare[3] customer reviews (written in English) about food-related POIs (i.e. POIs that belong to the food category) using the coordinates[4] of each city of interest. Overall, we created 10 collections (one for each city) as illustrated below in Table 1.

Table 1. Amount of data collected from Foursquare for each city.

City/country	Comments	POIs
Amsterdam/Netherlands	53.767	3.851
Athens/Greece	39.359	5.432
Barcelona/Spain	44.114	6.595
Berlin/Germany	29.386	6.998
Brussels/Belgium	14.872	5.428
Lisbon/Portugal	12.295	3.519
London/United Kingdom	158.284	16.888
Paris/France	32.382	17.250
Prague/Czech Republic	49.827	4.600
Rome/Italy	14.984	5.882
Total	449.270	76.443

2 http://stevianet.gr/dashboard.

3 https://foursquare.com/.

4 We used a specific bounding box for each city.

3.2 Topic Modelling

Topic models provide an effective way to obtain insights in large collections of unlabeled data and have been widely used for inferring low-dimensional representations that uncover the latent semantic structures of textual [8], image [9], or audio [10] data. A topic is a probability distribution over words, where distribution implies semantic coherence. The state of the art topic model method is Latent Dirichlet Allocation (LDA) [11] and its derivatives, whilst the recently proposed algorithms for Non Negative Matrix Factorization (NMF) perform also well for document clustering and topic modeling [12, 13].

We applied both LDA and NMF methods with a variable number of topics for each data collection. After many iterations and human evaluation, we decided to extract 50 topics for each collection. The results indicate that the NMF topics were of better quality in terms of topic covering and coherence, and by far faster to compute than the corresponding LDA ones. This is probably due to the tf-idf weighting scheme that the NMF is using. The input for the topic models was preprocessed data. In particular, all comments were lowercased and we removed the stopwords. Next we applied a part of speech tagger and after running some experiments we decided to keep only the nouns and the adjectives. Then, we extracted bigrams from each text targeting to obtain topics with phrase like keywords and not only single terms. For the NMF method we used tf-idf weighting scheme with l2 normalization to construct the terms – documents matrix. The output of the method is clusters of words indicating a topic (e.g. carbonara, amatriciana, ravioli, spaghetti, mimosa). No taxonomies or ontologies of topics were used to assist topic modelling. At a final step, a human evaluator inspected the output and assigned a descriptive label to the ten top topics for each city (e.g. PASTA for carbonara, amatriciana, etc. in the above example).

3.3 Sentiment Analysis

Sentiment Analysis can contribute to a better understanding of public opinions, emotions, needs and concerns. Hence, it is a key data analytics tool in the context of urban sensing and citizen behavior analysis used not only for textual data like Twitter [2, 14], but also for visual data shared on social networks (e.g. Pinterest) [15]. Sentiment Analysis solutions range from general purpose algorithms to more fine-grained approaches like aspect-based [16, 17] or topic-based [18] sentiment analysis depending on the case. State of the art techniques range from the traditional lexicon-based [19] to the current trend of deep neural network approaches [20].

We employ a neural network approach to build a model that classifies user reviews according to their positive/negative sentiment orientation. To train our model we used the Fine Foods Amazon Reviews Dataset [21] that consists of user reviews about Amazon food products. Each review is accompanied by a 1–5 score given by the users. Reviews with score 1 and 2 were labeled as Negative. For the Positive class we only use the 5-star reviews in order to create a more balanced training dataset, since the amount of the positive and negative reviews was disproportional in the original dataset. We then built a model that predicts whether a review is positive or negative. Our model uses pretrained word-embeddings; an embedding layer projects each token into an embedding space and each token is transformed into a vector. The vectors are initialized by pretrained GloVe vectors [22], but they can also change during the training process. Then a convolutional layer consumes the vectors and learns filters that read through the bigrams and trigrams of the text. We used a LSTM layer that reads through the representations of the bigrams and trigrams produced by the convolutional layer. In order to better match the vocabulary of the pretrained embeddings, the review texts were tokenized. The input for the model was collections of topic-specific user comments derived by the topic modelling analysis for each city. Each comment is classified as positive or negative with a confidence score between 0 and 1. If a comment is associated to more than one topics, the same sentiment value is assigned to all topics.

4 Results

4.1 Food Perceptions: Food-Related Topics and Sentiments

The topic analysis output is visualized using heat maps that provide a comprehensive overview of which geographical regions of a city are associated with each topic. For example, Fig. 2 portrays the spatial distribution of the topic "BURGER FOOD" in downtown Athens according to the spatial distribution of the Foursquare user-comments about food-related POIs. More intense color indicates areas with higher density (i.e. more comments for the particular topic in a specific area). The restaurant map icon () is used only for the POIs that are indexed as burger joints in Foursquare.

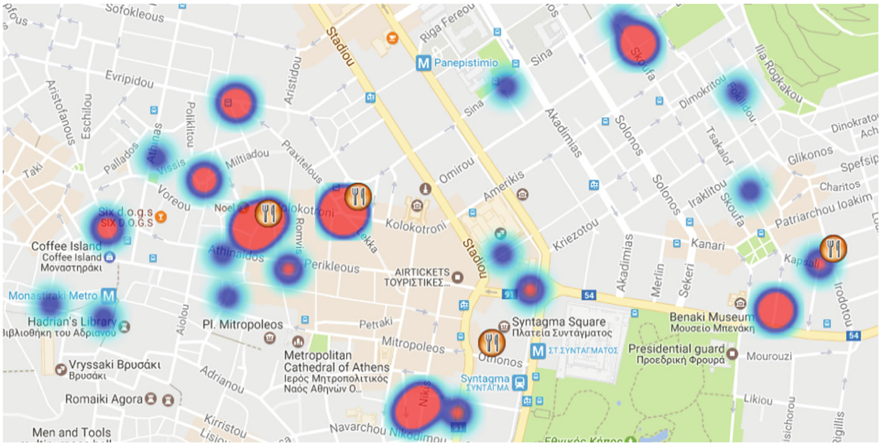

Fig. 2. Heatmap for the topic "BURGER FOOD" in Athens and POIs tagged as burger joints (Color figure online).

Given that the topics reflect the food consumers' and not the food providers' standpoint (i.e. owners of food-related POIs) they provide additional information about the food imprint of a city with regard to a specific topic. In other words, comments about burger food are not limited to burger joints. As it is illustrated in Fig. 2, there are much more burger food options in downtown Athens, indicating at the same time the mostly discussed, and by extension highly visited, ones. Integrating the sentiment analysis output (see Fig. 3) enables also answering to queries like "*Where can someone find good burgers/breakfast/pizza etc. in Athens/Rome/Paris etc.*", where "*good*" indicates positive user-comments about each topic.

Fig. 3. Heatmap of positive comments for the topic "BREAKFAST" close to Louvre Museum in Paris

The comparative analysis of the top ten topics for each city revealed that some topics are common in some cities, whilst others are unique per city indicating the local food identity/culture. For example, "ITALIAN CUISINE" is a common topic for 7 cities (Amsterdam, Athens, Berlin, Brussels, Lisbon, London, and Rome), whilst the more fine-grained topic "PASTA" is unique for Rome. Similarly, "BEER" appears in the top ten topics for the cities of Berlin, Brussels and Prague, that is cities, and subsequently countries, with brewing culture and history. As for the sentiment analysis output, "BEER" receives the most positive comments in Berlin. Further qualitative analysis of the results, could indicate also the reasons behind this.

Given that the topics and the sentiments result from Foursquare comments written in English they reflect mostly visitors' and not residents' perceptions and opinions. Hence, in order to be able to answer questions like "*Which city has the best beer*?" or "*Which is the best burger/breakfast in town*" in a more accurate and "comprehensive" manner, we need to take into consideration also the residents' perceptions and more data sources (e.g. geolocated Tweets written in other languages as well).

4.2 Further Insights

The topic modelling output includes also topics that have to do with the customers' experience in the restaurant such as the service, the ambience, the location and the value for money. Hence, if we also take into account the dimension of time (comments' timestamp), it is possible, for example, to monitor how the food prices are evaluated by restaurants customers in different time periods (e.g. tourist seasons) in different geographical areas of a city. The results can also be used to compare the particular restaurants' aspects across different cities and countries e.g. which city has the worst SERVICE according to users' comments? According to the sentiment analysis output, Berlin and Brussels receive the most negative comments. Further linguistic analysis could explain the reasons behind the positive or negative evaluations. For example, Fig. 4 below presents the distribution of the words "friendly" and "slow" in all user comments for each city. The least mentions of "friendly" and the most mentions of "slow" appear in the reviews for Berlin and Brussels, so perhaps this indicates that are not very satisfied with the staff attitude and efficiency in these two cities as compared to the other eight.

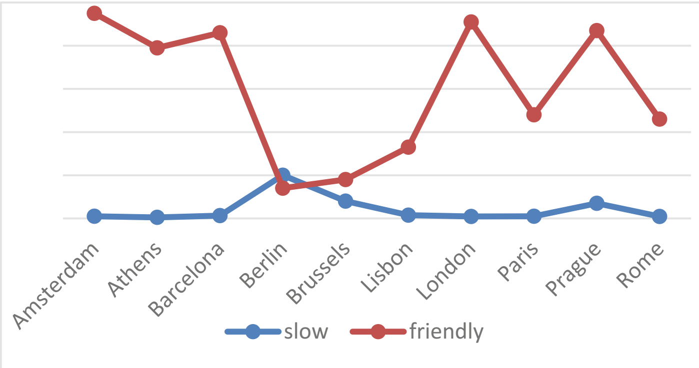

Fig. 4. Percentage of occurrences of the words "friendly" and "slow" in user reviews.

5 Conclusions and Further Work

In this paper, we presented an ongoing work exploring food experiences expressed in user reviews about food-related POIs in ten European cities with a focus on smart citizen and tourist sensing. The results indicate that the spatial distribution of food-related topics and sentiments enrich the cities maps with information that allows for a more insightful navigation across their different geographical regions, since they provide insights not available in the original data. Currently, we are working on further qualitative analysis of the results as well as on integrating results from other data sources (e.g. geolocated Tweets). Future work includes the extraction of other types of food-related insights and the combination of other types of geo-located data focusing on health issues such as obesity and diabetes, to investigate, for example, possible correlations between obesity in young ages and schools proximity to burger food.

Acknowledgements. Research presented in this paper is funded by the CAPSELLA (688813—CAPSELLA—H2020-ICT-2015) project.

References

1. Tajbakhsh, M.S., Bagherzadeh, J.: Microblogging hash tag recommendation system based on semantic TF-IDF: Twitter use case. In: 2016 IEEE 4th International Conference on Future Internet of Things and Cloud Workshops (FiCloudW) (2016)
2. Villena-Román, J., Luna Cobos, A., González J.C..: TweetAlert: semantic analytics in social networks for citizen opinion mining in the city of the future. In: UMAP Workshops (2014)
3. Elmongui, H.G., Mansour, R., Morsy, H., Khater, S., El-Sharkasy, A., Ibrahim, R.: TRUPI: Twitter recommendation based on users' personal interests. In: Gelbukh, A. (ed.) CICLing 2015. LNCS, vol. 9042, pp. 272–284. Springer, Cham (2015). doi:10.1007/978-3-319-18117-2_20
4. Karidi, D.P., Stavrakas, Y., Vassiliou, Y.: A personalized Tweet recommendation approach based on concept graphs. In: 13th IEEE International Conference on Ubiquitous Intelligence and Computing (UIC) (2016)
5. Joseph, K., Tan, C., Carley, K.: Beyond "local," "categories" and "friends": clustering foursquare users with latent "topics". In: Proceedings of ACM Ubicomp 2012, pp. 919–926 (2012)
6. Cheng, Z., Caverlee, J., Lee, K., Sui, D.Z.: Exploring millions of footprints in location sharing services. In: ICWSM 2011 (2011)
7. Yang, D., Zhang, D., Yu, Z., Wang, Z.: A sentiment-enhanced personalized location recommendation system. In: 24th ACM Conference on Hypertext and Social Media (2013)
8. Mcauliffe, J.D., Blei, D.M.: Supervised topic models. In: Advances in Neural Information Processing Systems, pp. 121–128 (2008)
9. Chong, W., Blei, D., Li, F.F.: Simultaneous image classification and annotation. In: Computer Vision and Pattern Recognition (2009)
10. Hoffman, M., Cook, P., Blei, D.: Bayesian spectral matching: turning young MC into MC hammer via MCMC sampling. In: International Computer Music Conference (2009)
11. Blei, D.M., Ng, A.Y., Jordan, M.I.: Latent dirichlet allocation. J. Mach. Learn. Res. **3**, 993–1022 (2003)
12. Kuang, D., Choo, J., Park, H.: Nonnegative matrix factorization for interactive topic modeling and document clustering. In: Celebi, M.Emre (ed.) Partitional Clustering Algorithms, pp. 215–243. Springer, Cham (2015). doi:10.1007/978-3-319-09259-1_7
13. Greene, D., Cross, J.P.: Exploring the political agenda of the European parliament using a dynamic topic modeling approach. Polit. Anal. **25**(1), 77–94 (2017)
14. Estévez-Ortiz, F.-J., García-Jiménez, A., Glösekötter, P.: An application of people's sentiment from social media to smart cities. El profesional de la información **25**(6), 851–858 (2016)
15. Ahmed, K.B., Bouhorma, M., Ahmed, M.B.: Smart citizen sensing: a proposed computational system with visual sentiment analysis and big data architecture. Int. J. Comput. Appl. **152**(6), 20–27 (2016)
16. Liu, B.: Sentiment Analysis and Opinion Mining. Morgan & Claypool Publishers, San Rafael (2012)
17. Pontiki, M., Galanis, D., Papageorgiou, H., Manandhar, S., Androutsopoulos, I.: Semeval-2015 task 12: aspect based sentiment analysis. In: 9th SemEval (2015)
18. Thelwall, M., Buckley, K.: Topic-based sentiment analysis for the social web: the role of mood and issue-related words. J. Am. Soc. Inf. Sci. Technol. **64**(8), 1608–1617 (2013)
19. Taboada, M., Brooke, J., Tofiloski, M., Voll, K., Stede, M.: Lexicon-based methods for sentiment analysis. Comput. Linguist. **37**(2), 267–307 (2011)

20. Socher, R., Perelygin, A., Wu, J., Chuang, J., Manning, C.D., Ng, A.Y., Potts, C.: Recursive deep models for semantic compositionality over a sentiment treebank. In: 2013 Conference on Empirical Methods in Natural Language Processing, pp. 1631–1642 (2013)
21. McAuley, J., Leskovec, J.: From amateurs to connoisseurs: modeling the evolution of user expertise through online reviews. In: WWW (2013)
22. Pennington, J., Socher, R., Manning, C.M.: GloVe: global vectors for word representation (2014)

Improvement of the Workers' Satisfaction and Collaborative Spirit Through Gamification

Evdoxia Eirini Lithoxoidou, Stefanos Doumpoulakis, Athanasios Tsakiris, Stelios Krinidis(✉), Dimosthenis Ioannidis, Konstantinos Votis, and Dimitrios Tzovaras

Centre for Research and Technology Hellas/Information Technologies Institute (CERTH/ITI), 57001 Thessaloniki, Greece
{elithoxo, stefanos.d, atsakir, krinidis, djoannid, kvotis, dimitrios.tzovaras}@iti.gr

Abstract. Supporting the use of technology into industrial environments is an issue of mass appeal within the Industry 4.0 initiative, with a lot of promising research especially into interaction, production and training sections. In this paper, a Social Collaboration platform is introduced which creates an online community for workers as well as an Augmented Reality (AR) tool for production and training purposes both of which constitute gamified applications on which a customizable Gamification platform is applied according to impending needs. These tools have been implemented in order to become of daily use to employees in factories and incentivize them to promote collaboration, engagement, participation and work satisfaction. Every promoted industrial behavior can be described and awarded based on the offered rule engine. Thus, two gamified processes are offered regarding Social Collaboration and AR training to employees of industrial environments, so that an inference is drawn concerning participation, satisfaction and self-fulfillment. The use of the platforms is illustrated in this paper by two examples consisting one use case for a section in social collaboration and a second use case of training through AR.

Keywords: Gamification · Social collaboration · AR training · Industrial environment · Factories of the future

1 Introduction

In an industrial environment, jobs tend to be repetitive and sometimes tedious so it is really important to create challenges for engaging workers. Besides, as long as online communities are evolving continuously, great opportunities are offered for connecting and interacting with colleagues on a daily basis. Many tools are implemented in order to support collaboration between workers inside the factory and Social Collaboration Software is one of them. It can be used to exchange knowledge but also social content through sound, video or text. It consists of a vast offering of platforms that concern enterprises [1], software team collaboration [2] and in general the positive impact of social software solutions in the acceleration of problem solving and in increasing the user's interaction with the system [3]. In general, social collaboration software is a way

I. Kompatsiaris et al. (Eds.): INSCI 2017, LNCS 10673, pp. 184–191, 2017.
https://doi.org/10.1007/978-3-319-70284-1_15

of collecting data to support human collaboration within groups and organizations [4]. Apart from Social Collaboration platforms, Augmented Reality (AR) training presents an emerging domain of improving technical skills and enriching collaboration. There are platforms that use augmented reality techniques for educational purposes [5], for healthcare [6] as well as for engaging elderly or impaired persons in production [7]. Especially in industrial environments where maintenance and assembly tasks are of high importance, AR training is a powerful technology [8] that can be supported by gamification too.

In order to create successful gamified processes there are three main principles that should be fulfilled [9]. These are relatedness, competence and autonomy for motivating people intrinsically while completing the "self-determination theory" described by Deci and Ryan [10]. After having studied about how to engage people into actions and boost their interest and satisfaction, research has shown that this is feasible by building self-confidence in a pleasant environment.

The remainder of this paper is structured as follows. In Sect. 2 the design of the gamification engine is described as well as the route that a gamified action is following until fulfillment. Section 3 introduces two use cases on which gamification elements are applied and in Sect. 4 there is a short description of how the engine supports both of them. In Sect. 5 a link is provided for demonstration and finally in Sect. 6 there is reference to future work along with conclusion.

2 Architectural Structure and Business Logic

The gamification platform implemented is based on a server-client architecture which is a computing model that delivers and manages web services that are consumed by the client. The Gamification engine can have many instances and each one of them stores data in a database that is connected with the server. It provides the ability to create a game either for team playing or individual playing [Fig. 1]. In order to use the gamification engine for creating gamified processes, one has to be the administrator of the platform. Users of the platform can be of two general categories, the administrators and the players where according to the nature of the workplace administrators are managers, supervisors or trainers and players are technicians, operators and employees. Their role is discrete, both categories can participate in the gamified processes, but administrators can set the rules of the processes as well. This interface provides completely customizable features where the administrator manages the whole game's behavior by setting rules and awards of four types. The award types are distinguished between scalar, sets, tangible objects and levels and they can be viewed on leaderboards that concern either individuals or teams. Scalar awards can be points or coins while sets are groups of badges including cups, medals or anything that the administrator of the platform can imagine. Additionally, when talking about tangible awards, they are also groups of material or monetary incentives or maybe incentives that can easily be translated into a cash value like coupons. Finally, levels describe the experience of a player in the game by giving titles like "Novice", "Intermediate", "Master" or anything similar.

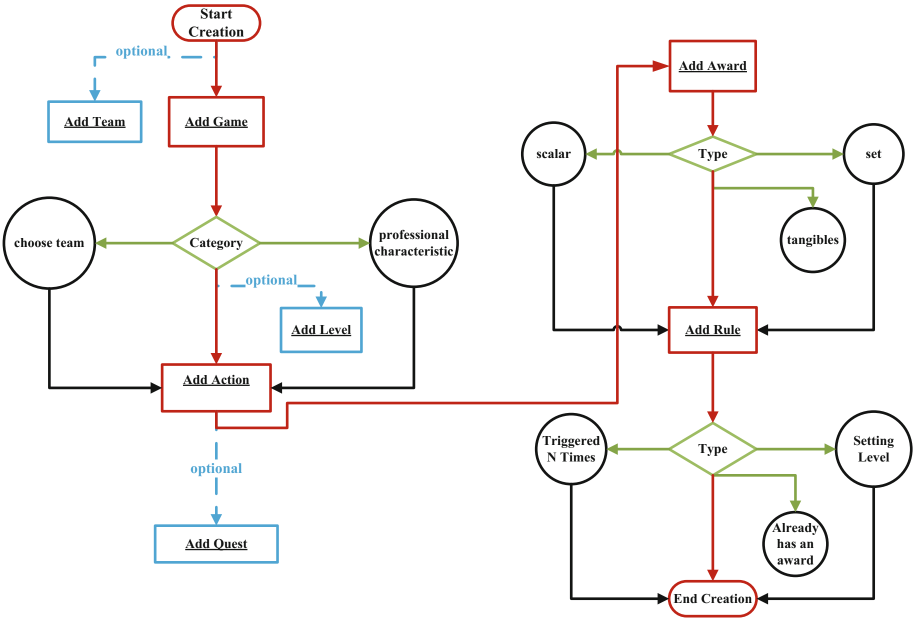

Fig. 1. Gamify a process

The creation of a gamified scenario begins with administrator deciding which process to gamify and choosing if it is going to be addressed to individuals or to teams. Processes already exist, so the players do not have to include any extra activity in their workday. Job is fulfilled in the same manner as before and awards are gained. After defining the kind of game (individual or team game) the individual's category or the team is chosen as well. In order to gamify work activities, actions are created and associated with them. Sequence of actions form a mission in a custom game. Besides, quests are actions which can be fulfilled only within a time interval thus consider the action "Ask Question" which if fulfilled for a specific amount of times and in between the defined time interval, the player gains an extra award. Usually these activities are presented in notifications, informing the user of their availability.

Every game has its rules, thus, the idea is to think about the relations between actions and awards and capture them by the gamification platform's rule engine. Every action that is triggered follows the chart presented in Fig. 2.

The administrator can combine all these elements by using the rule engine found in the Rules section. Rules are the basis of the gamification process because they define the behavior that a user must follow in order to achieve the award. There are three kinds of rules that describe the following behaviors:

1. Rule for N times triggered: This type of rule determines what the player gains after fulfilling an action from 1 to N times.
2. Rule for having the award: After having collected some scalar awards, the platform gives the ability to exchange them with badges or tangible awards. This second type

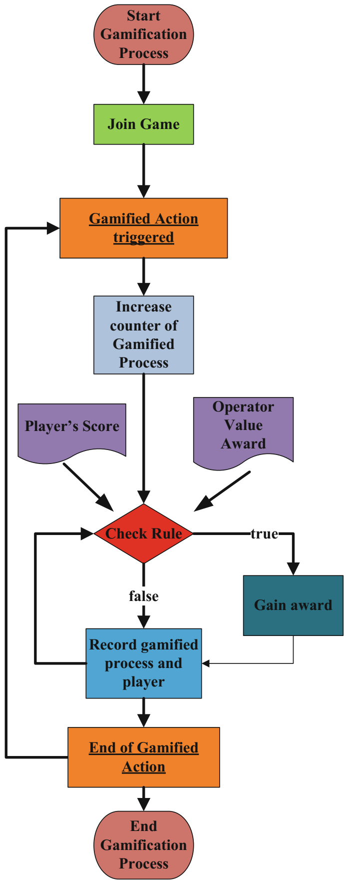

Fig. 2. Gamified applications' flow

of rule checks the amount of scalar awards gained from a user and accordingly awards the player with the equivalent set award.

3. Rule for setting level: In order to advance levels, the engine checks the amount of scalar awards collected and when they are equal to or more than the starting point of a level, the player advances to it. For levels, there is always a hint to reveal a way for gaining it which is provided from the administrator as well, in order to drive behavior. Games do not have fixed labels in order to improvise with their labels.

Moreover, there is a leaderboard in every gamified process which informs the administrator of the ranking and provides key insights when combined with what it actually means regarding worker performance and participation.

Every user can see the available games in the profile tab, in achievements, and join or leave a game according to preference. Each game provides a leaderboard, in order for the user to know who is already participating and how many points are already gained before joining it [Fig. 3].

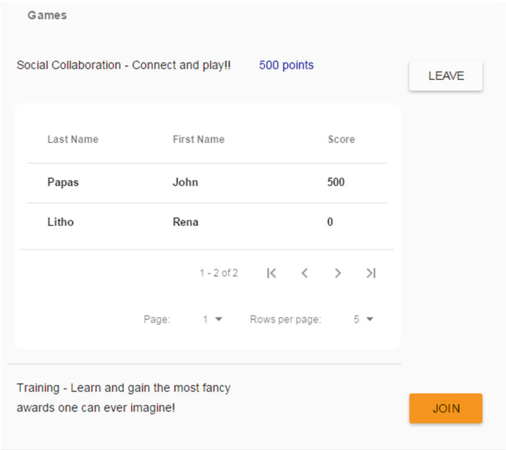

Fig. 3. Available games and participants

Every award either of scalar, set or level type is shown in the profile's tab, to every user separately and personally while every user can see her connections' achievements too.

In this system the concept is to provide a flexible way of creating awards according to each desired application. In other words, it is important for the administrator to know the workers' needs and provide appropriate awards as motivation to fulfill tasks. In this platform there is no restriction in setting an award of any kind as the engine is able to recognize it.

The Gamification platform can be connected with every tool that can gamify its actions via API calls to the web services that are implemented within this scope. It offers a multilingual interface for the end users by providing the ability of selecting from a set of supported languages when creating the game. The administrator of the platform is offered English as default and can add an accurate translation of the supported languages. Furthermore, the gamification platform keeps historical logs of all actions triggered from every user and every game that has been created, along with the awards that are gained, in order to provide an informative data analysis.

The platform's concept is to be an analytical platform for addressing demands of users, measurements and inference. For this goal, the administrator of the platform needs to know the list of behaviors needed to be tracked in order to recognize which sections need improvements and how. Once the gamified processes are set, users can start playing so the gamification platform needs to take care of unintended behavior. There are rules which define the maximum number of every award that can be gained as a limit to how many times one can fulfill an action. These boundaries can mitigate the possibility of people overdoing it with specific behaviors.

3 Gamified Training and Productivity Applications

3.1 Social Collaboration Platform

This feature encourages the exchange of knowledge and support to each other. The Social Collaboration platform provides the possibility to the users to connect with each other and ask or seek for help when needed. Additionally, they can support each other through this support-oriented environment where questions can be asked and answers can be given by the users. Novice workers receive support by the experienced ones and this leads into reducing stress, enhancing cooperation and avoiding mistakes in a safe support-oriented environment. Every interaction with this platform, when it aims to cooperation among users is gamified and rewarded. There are rules that evaluate the fulfillment of actions, the frequency of the event and the time interval needed while using filters that regard different teams, ages and departments.

3.2 AR Training

The application of gamification on the AR Training Platform provides incentives for workers to follow the training regime more diligently. At the same time it provides trainers with evaluation methods of training performance for each trainee as well as cumulative insights regarding the training effectiveness as a whole, through time, per different age group etc. Actions performed on the AR Training platform are matched to equivalent gamified activitied on a dedicated "AR Training Game" where points, achievements and levels correspond to increased performance and participation in the training sessions.

4 Implementation and Use

4.1 Gamification Support to Social Collaboration

Gamification techniques have been applied on a Questions and Answers forum section of Social Collaboration platform. The need arose in industrial places with many employees, who socialize, communicate and share same issues when gathered together but are not able to do this anytime in a flexible way. Therefore, the Social Collaboration platform offers a digital gathering environment for fostering employees' interaction and collaboration. In order to motivate them, Social Collaboration actions have been gamified by the administrator of the platform according to their needs.

After fulfilling gamified actions from Social Collaboration, the user can see the gained points and badges as well as the achieved level in the personal profile [Fig. 4].

The platform has been deployed in three pilot sites in order to acquire feedback and evaluate responsiveness, engagement and satisfaction of participation.

4.2 Gamification Support to AR Training

In the same manner as in the previous use case, the AR Training game consists of rules applied to different aspects that determine training evaluation. These aspects include the

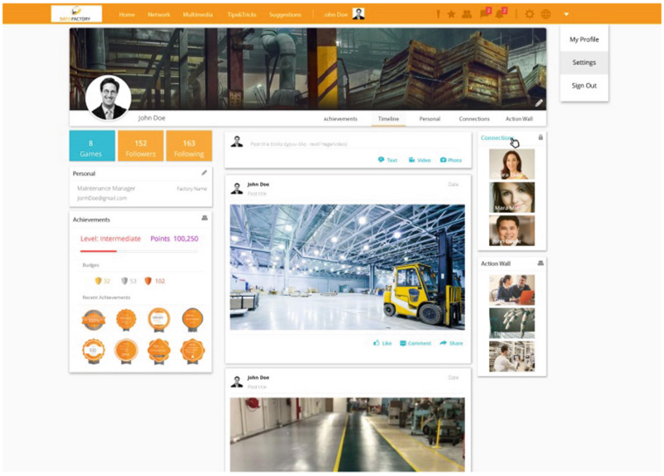

Fig. 4. Profile page for social collaboration with gamification elements

number of steps completed successfully, overall training session duration, when applicable, the number of errors performed etc. These rules can be customized by applying specific values to the awards, points, levelling etc. during the authoring of a training session to specific training activities within each session.

Gamification feedback is both direct within the training app itself upon completion of each training session as well as from within the profile page of each user in the Social Collaboration platform. Each time a worker performs a training session using the AR Training Platform Presentation tool, points are added to the game for each step correctly performed. Furthermore, awards are given when multiple training sessions are completed and over time, based on the performance of the workers on the sessions and the points gathered, they advance on a levelling system.

The AR training platform has been deployed in three pilot sites in order to acquire feedback and evaluate training performance using novel tools, improvement of expertise and faster training output. This is complemented with the connection of the AR Training platform to the Gamification Platform and the Social Collaboration platform.

5 Demonstration

The described gamification engine as well as the two use cases are demonstrated on a video uploaded to Google Drive:

https://drive.google.com/open?id=0B0l69iMs7J4DZXR5VUk2NTdXbG8.
For better quality one should download it.

6 Conclusion and Future Work

This work is implemented in order to provide a complete system of gamification tools which can be applied on a large amount of use cases. Functionality is based on web services which can be invoked from any application that connects to the API. The fact that the platform keeps a historical database provides the ability to have data analytics in order to understand the behaviors that are driven in fine-grained details. Thus, as mentioned above, there are three pilots that are going to use the platform described in this paper so that an analysis could verify if engagement to tasks has been improved after including gamification. In the future, a statistical analysis is going to be performed, originated from actual daily use which is going to help improve the key performance indicators (KPIs). Afterwards, according to this analysis, the first version of rules of the gamified processes is going to be updated in order to oblige the ongoing needs.

Acknowledgments. This work has been partially supported by the European Commission through the project HORIZON 2020-INNOVATION ACTIONS (IA)-636302-SATISFACTORY, and the project HORIZON 2020-RESEARCH & INNOVATION ACTIONS (RIA)-690211-CAREGIVERSPRO.

References

1. Li, M., Chen, G., Zhang, Z., Fu, Y.: A social collaboration platform for enterprise social networking. In: IEEE 16th International Conference on Computer Supported Cooperative Work in Design 2012 (CSCWD), pp. 671–677 (2012)
2. Wolf, T., Schröter, A., Damian, D., Panjer, L.D., Nguyen, T.H.: Mining task-based social networks to explore collaboration in software teams. IEEE Softw. **26**(1), 58–66 (2009)
3. Stieglitz, S., Lattemann, C., Robra-Bissantz, S., Zarnekow, R., Brockmann, T. (eds.): Gamification: Using Game Elements in Serious Contexts. Springer, Cham (2017). doi:10.1007/978-3-319-45557-0
4. Iyer, L., Paul, S., Zhou, L.: Introduction to the data science for collaboration minitrack. In: HICSS, p. 337 (2016)
5. Brown, T.M., Smith, T.R., Gabbard, J.L., Gilbert, J.E.: Advanced learning technologies (ICALT). In: 16th International Conference on IEEE 2016, pp. 260–264.IEEE (2016)
6. Pereira, P., Duarte, E., Rebelo, F., Noriega, P.: A review of gamification for health-related contexts. In: Marcus, A. (ed.) DUXU 2014. LNCS, vol. 8518, pp. 742–753. Springer, Cham (2014). doi:10.1007/978-3-319-07626-3_70
7. Korn, O.: Industrial playgrounds: how gamification helps to enrich work for elderly or impaired persons in production. In: Proceedings of the 4th ACM SIGCHI symposium on Engineering interactive computing systems, pp. 313–316 (2012)
8. Webel, S., Bockholt, U., Engelke, T., Gavish, N., Olbrich, M., Preusche, C.: An augmented reality training platform for assembly and maintenance skills. Rob. Auton. Syst. **61**(4), 398–403 (2013)
9. Groh, F.: Gamification: state of the art definition and utilization. Inst. Media Inform. Ulm Univ. **39**, 39–46 (2012)
10. Deci, E.E., Ryan, R.M.: Intrinsic Motivation and Self-Determination in Human Behavior. Springer, US (1985)

Online Grocery Shopping: Identifying Change in Consumption Practices

Jo Munson(✉), Thanassis Tiropanis, and Michelle Lowe

Web Science Institute, University of Southampton, Southampton, UK
{J.Munson,T.Tiropanis,M.Lowe}@soton.ac.uk

Abstract. Following the invention and proliferation of the Internet, Web and mobile technologies, we have seen a global revolution in retailing. Despite the rapid growth of e-commerce, the online grocery shopping market has taken until now to gain traction, currently constituting 6.9% of the UK's grocery market, but projected to increase 68.3% to £17.6 bn by 2021. There is little work accounting for new and contingent behaviours in the online grocery market, not least because of historically poor access to retailers' data. This paper leverages access to the UK's fourth largest supermarket, WM Morrisons Plc (Morrisons) to investigate consumer behaviour in this market, augmenting the Office for National Statistics' Living Costs and Food Survey, the UK's only substantial publicly available resource to date. This paper establishes that there have been changes in consumer behaviours in response to the unique opportunities and challenges of online grocery shopping and explores the specific socio-technical factors that may be contributing to these changes, namely: ease of price comparison; attitudes to purchasing perishable goods online; and logistical considerations. Furthermore, it provides some evidence that the proportion of fresh products bought online exceeds the proportion bought offline, contrary to popular belief. Finally, this paper argues that with correction for location bias, the Morrisons sample could provide a proxy for examining online grocery behaviour in-depth at the national level.

Keywords: Online grocery shopping · E-commerce · Consumer behaviour · Retail analytics · Socio-technical systems

1 Introduction

Following the invention and proliferation of the Internet, Web and mobile technologies, we have seen a global revolution in retailing. Online retail sales are surpassing expectations in the UK, exceeding £133 bn in 2016, a 16% increase on 2015 [3]. As at March 2017, the Office for National Statistics (ONS) report that 15.0% [15] of all retail spending now occurs online (up from 8.5% in 2012 [34].) Despite this surge in online shopping, one sector that has been slow to embrace the e-tail revolution is grocery shopping. The slower growth in this sector has been attributed to two factors: supply-side issues (such as low margins and the

I. Kompatsiaris et al. (Eds.): INSCI 2017, LNCS 10673, pp. 192–211, 2017.
https://doi.org/10.1007/978-3-319-70284-1_16

logistical constraints of low-density deliveries [22,31]); and conservative demand. Low demand is often attributed primarily to consumers' desire to pick their own perishable goods such as fresh fruit, vegetables and meat [24,43]. Despite significantly slower market penetration, the UK is one of the few countries that has established a growing online grocery market, currently worth c.£10.5 bn p.a. [21]. Online food shopping currently constitutes 6.9% of the UK's grocery market [48] and IGD predict the value of the UK's online grocery market will increase 68.3% to £17.6 bn by 2021 [21]. Tesco claim to have made profits of £127 m from online grocery sales in 2013 [32], whilst Pure-play retailer Ocado posted profits of £14.5 m over 2016 [6].

With growth now well established in the UK's online grocery market, it is a poignant time to investigate whether there have been changes in grocery consumption behaviour in the new online context. As in other e-tail sectors, a number of market analysts have argued that there is a huge opportunity in online grocery shopping to gain insight into consumer behaviour via "Big Data" accumulated by online consumption [27,33]. Despite this, there is little work accounting for new and contingent behaviours, not least because of historically poor access to retailers' data. Understanding behavioural change entails game-changing potential for retailers, particularly in a low-margin industry such as the UK grocery market. The opportunities afforded by increased insight are numerous with respect to customer retention, complementary service provision, online and physical store planning, better lifetime value prediction, personalisation and brand reputation. The usefulness of understanding online grocery consumption is not limited to retailers however. Governments and town planners can benefit from understanding movement flows and practices in order to plan services more effectively. Technology designers can benefit from understanding technology use, especially in the interpretive flexibility of technologies. From an academic perspective, insights into consumer movements, habits and shopping practices contribute to the broader thesis of understanding socio-technical systems, social change and consumption behaviour.

To date, the most comprehensive publicly available data covering online grocery shopping in the UK is published by ONS in their annual Living Costs and Food Survey (LCF), the latest of which separates online grocery shopping from supermarket and other food outlet shopping [5]. Whilst the LCF provides a valuable starting point for uncovering changes in grocery consumption behaviour, its scope is limited to spending by food category for a sample of c.4,000 households. The LCF does not collect and collate information about use of mobile devices, how users interact with online stores or where and when they do their grocery shopping. The authors of this paper have been granted access to a large and high quality dataset of real online grocery transactions via the Google Analytics account of WM Morrisons Plc (Morrisons). Google Analytics allows analysis of consumer activities at the level of the individual and facilitates analysis of the place, device, timing and browsing behaviour of customers.

Morrisons does not currently deliver nationally and the customer base is significantly overrepresented in the North West and Yorkshire and the Humber regions relative to the general population. As a corrective to the potential bias in the sample we re-weight findings using population averages and show that there is significant potential to use the Morrisons sample to explore online grocery consumption at the national level. This paper establishes that that there have been changes in consumer behaviours in response to the unique opportunities and challenges of online grocery shopping and explores the specific socio-technical factors that may be contributing to these changes.

1.1 Aims of the Study

This paper provides an initial assessment of the potential to use the Morrisons dataset for exploring national level online grocery behaviour and examines in detail the features of online grocery that are thought to constrain or change behaviour - customer reticence buying perishable and tangible food stuffs; the ability to compare prices with Web search; and the perceived tendency for consumers to revert to "favourites" or previous orders. Specifically, this paper addresses the following questions:

– Is the Morrisons sample representative of online grocery shopping at the National Level?
– Are there differences in consumption behaviours in online and offline grocery shopping in the UK?
– Are online grocery shoppers price sensitive?
– Are online grocery baskets stable?

The next section of this paper gives a brief overview of relevant grocery and e-commerce research and identifies gaps in the literature motivating this work.

2 Background

2.1 Online Grocery Shopping - Adoption, Retention and Motivation

Attracting and retaining customers has always been of primary concern to retailers and with some estimating that a 5% increase in customer retention can lead to profit increases of nearly 100%, it is easy to see why [38]. Rafiq and Fulford [37] found that offline interaction with the brand and "word of mouth" recommendation online were the most important predictors of choosing and continuing to use the services of an online retailer. Rafiq and Fulford [37] also found that market leader Tesco's success in gaining customers from other retailers was an exemplification of the 'double jeopardy effect', i.e. that those with lower market share also suffered lower brand loyalty. This effect was also shown to be true for individual products in Danaher et al.'s study of grocery shoppers in New Zealand [11]. Dawes and Nenycz-Thiel analysis of Kantar Panel data indicated that loyalty to a single retailer was less (and falling) online, but that brand loyalty (including retailer own-brand products) was slightly higher online [12]. This may relate to brand portrayal online or the stabilising effects of "favourites" baskets, addressed in this paper.

HCI, psychology and market research analysis of online consumption has centred largely on brand-impervious motivations for technology adoption. Hansen [19] postulates that OGS is a 'discontinuous innovation' requiring a significant shift in behaviour and thus a more drawn-out adoption period. Hand et al. found that being adept internet users was not sufficient to assume a propensity to engage with online shopping and that when online shopping was adopted, it did not usually entail the discontinuation of offline shopping [18]. Several studies have employed various incarnations of the Technology Acceptance Model (TAM). In one such study, Evanschitzky et al. [17] recognised the importance of context in the adoption of technologies. They identified the 'human-to-human' consumer-retailer interaction as having been replaced with 'human-to-computer' interaction as a key differentiator between offline and online customer satisfaction. They reported a moderately good fit with Szymanski and Heise's results concluding that some drivers of E-satisfaction may be context invariant [17,42]. TAM studies tend to be deployed in experimental settings and often involve the "simulation" or "intention" to adopt a technology or behaviour, but do not assess the uptake or continued use of the technology. This is justified on the assumption of an 'Intention-Behaviour' link, i.e. the assumption that a customer's reported intentions are played out in their subsequent actions. Bagozzi describes the link as 'probably the most uncritically accepted assumption in social science research' [2]. Bagozzi goes on to criticise TAM for its simplistic model of human action, which fails to recognise that engagement with a given technology or service is rarely a "terminal goal" in itself for the consumer. Instead, he suggests that intentions and goals are continuously negotiated as users look to overcome obstacles, resist temptations and maintain willpower to achieve evolving goals [2].

Whilst these studies provide a good exploratory starting point for understanding consumption behaviour, they have rarely been verified or explored in combination with real-world online transaction data. This study uses an extensive corpus of real-world data to assess whether there are differences in consumption behaviour between online and offline grocery shopping in the UK.

2.2 Mobile Technologies and Online Grocery Shopping

As we move into an era where technological interfaces form the basis of many of our social and commercial activities, consumer behaviour has acquired a further layer of complexity [23]. Increased reliance on the Web presents potentially transformative effects in terms of our individual relationships with food, shopping and roles within the household; but also raises questions about the design and effects of the interfaces themselves. Mobile technologies not only pervade our homes, but also move with us, facilitating a relocation of (grocery) consumption from the physical store to "potentially anywhere". Crewe and Lowe [10] discussed spatial heterogeneity in terms of 'micro-geographies of consumption' where they argued that retailers create highly individual consumption spaces, quite apart from the perceived globalisation effect [39]. The implications of consumption that happens "anywhere" entails that consumption depends not only on the virtual space that is the Web interface, but also on the diverse array of physical spaces in which the consumption takes place. In focus group discussions with online shoppers, Michaud Trevinal and Stenger [28] found that many respondents enjoyed the affordances of being in a familiar, relaxing environment with the capacity to move between tasks, or research items on different websites when shopping online. The appeal of this ability to "multi-task" was echoed by respondents in Robinson et al.'s study [40]. Robinson et al. [40] conducted focus groups with 32 online grocery shoppers and found that respondents rarely purported to having a regular time or place for shopping, 'sometimes from home, sometimes from work, other times in the middle of the night.' [40].

Maity and Dass [25] found that the channel of consumption affected decision making and propensity to buy. Respondents preferred shopping in-store to on mobile devices, which the authors attributed to mobile phones' low 'media richness'. The convenience of e-commerce (via a desktop/laptop device) surpassed its relatively lower media richness compared to in-store, rendering it the most popular channel. Wang et al.'s study of M-shopping indicated a tendency to opt for branded or "known" products due to the constraints of the small screen size making it harder to research and evaluate new items [47]. Briesch et al.'s [4] offline study of 11,005 store visits indicated that retailers with more brands increased their probability of a household choosing their store, but that those with fewer SKUs per brand also attracted greater loyalty. Wang et al.'s 2015 study of c 16,000 consumer activities on a US-based online grocery platform represents the most comprehensive study to date [47]. The study used real-world data to examine the effect of mobile use on consumption behaviour. They found that consumers using mobile devices for online grocery shopping, 'M-shopping', increase in value over time, ordering more frequently as they become accustomed to the technology and interface. They also suggested that engagement with M-shopping resulted in low income consumers spending more than they did prior to commencing M-shopping. We are not aware of any study of this scale looking at UK audiences.

With notable exceptions [4,12,47] most online grocery shopping research to date has been qualitative in nature or based on self-reporting of behaviour and preferences in small-scale studies. This study redresses this balance by considering the real-world behaviours of hundreds of thousands of consumers and looks specifically to identify whether there are changes in consumption activities in the new online context. Future work will also look at identifying and verifying changes in consumption in the wake of the mobile technologies.

2.3 Trust, Price Sensitivity and Thrift

Another common theme emanating from previous studies is the lack of trust associated with purchasing fresh produce online. It has been found that consumers are worried about substitutions, the shelf-life of perishable goods and sub-optimal selection of fresh produce [18]. As a result, it has been reported that many online grocery shoppers continue to visit physical stores to purchase fresh food items [18]. This could have implications for all food stuffs, but particularly "fresh", where customers usually rely on sensory perception to evaluate items. This may imply a tendency to opt for branded, pre-packaged fresh goods. This study will add to this debate by comparing online and offline consumption of perishable and nonperishable goods.

Grocery shopping literature has long focussed on consumers' ability exhibit good economic sense, or "thrift". In his 1998 book, Miller claimed, 'That which the shopper does on behalf of the household is governed by thrift, while their individual presence is signified by the treat.' [29]. If this observation has merit, then one might suppose that the asocial nature of shopping on a personal device might lead to a higher propensity to "treat" oneself. Anecdotal evidence [16] suggests that this is not the case - that in fact online grocery shopping is used as a tool for planned, thrifty shopping and that it is the physical convenience stores that are frequented to top up on the treats that one tried to abstain from in the weekly shop. Robinson et al. [40] also alluded to a potential shift in shopping consumption behaviour associated with showing thrift in the online context. They found that users place orders less regularly to justify the cost of delivery charges, but that baskets may also be correspondingly bigger, as to diminish the marginal cost of delivery. Despite this, Huang and Oppowal found that physical distance from the supermarket was a more important predictor of tendency to shop online than delivery charges [20]. This was consistent with the findings of Briesch et al. [4], who found distance to travel has a larger effect on retailer selection than price or product assortment. In the online context however, where distance is removed as a variable among online retailers, price and product variety may play a larger role in choice of retailer. Contrary to retailer's fears that price comparability online would spark intensified price competition, Degeratu et al. found that online consumers may not be as price sensitive as the general population. They also reported that brand loyalty was more likely to persist online where there was a paucity of information [13]. It should though be noted that the online and offline groups were distinct. Despite attempts to choose similar demographics across the two groups, this does raise the question as to whether variation between channels could merely represent differences inherent in the online and offline groups [7].

This study contributes to this debate by addressing the question, 'Are Online Grocery Shoppers Price Sensitive?'. This will be achieved by looking for evidence of thrifty behaviours in site navigation and by comparing the average size of online baskets with offline national averages.

2.4 Time Poverty and Basket Stability

Twenty-first century families, particularly "working mothers" are often referred to as 'time poor' Wajcman [46] and time is often cited as the primary reason for online shopping. This was supported by the findings of Anesbury et al. [1], who reported that online shoppers spent a matter of seconds selecting each product, akin to offline findings [8,14] and concluded that there was little evidence of a change in behaviour between online and offline shopping in respect of time. They also noted that very few customers made use of the customisation options for displaying products, preferring to use the "search bar" and default product layouts.

This study did however rely on artificial shopping environment where consumers were asked to purchase a list of items and did not have to transact/pay for the shop, thus undermining the realism of their shopping behaviour. Time poverty has not been universally reported in studies to date. Whilst Robinson et al. [40] found that regular shoppers had reduced their shop time to minutes, facilitated by features such as “favourites”, Michaud Trevinal and Stenger [28] described how users’ propensity to multi-task meant that online shopping was done over an extended period. Furthermore, Huang and Oppowal found no evidence to support time being a factor of convenience for online over physical shopping [20]. Rohm and Swaminathan’s [41] online shopper typology cites four distinct user types and may help account for disparity in findings regarding time ad price sensitivity. Rohm and Swaminathan’s “convenience” motivated group were the most likely to engage in online shopping [41]. Their characteristic attributes included lower requirement for variety (across retailers) and lower sensitivity to receiving products immediately than the supermarket shopping population at large. The use of “favourites” and site search will be scrutinised in this study, when looking to establish whether online grocery baskets are quick to stabilise.

This section has summarised literature in the field of online grocery shopping and has identified gaps in the literature, namely a paucity of empirical real-world quantitative analysis of online grocery shopping; and little empirical evidence of user’s response to the new online context in terms of price sensitivity, total spend and product selection. In this paper, we focus on establishing whether there is evidence of differences between online and offline grocery consumption. We also examine whether online grocery shoppers are price sensitive and difficult to upsell or cross-sell to.

3 Methodology

This study employed an exploratory approach using Morrisons’ online transaction data and comparing it with online and offline national average estimates, as produced by the Office for National Statistics (ONS) in the Living Costs and Food Survey 2016 (LCF 2016) [5]. In so doing, this study examined the extent to which the Morrisons sample can be used to investigate online grocery shopping at the national level.

To establish whether the composition of Morrisons’ online baskets was comparable with national estimates for online and offline baskets, an analysis of basket composition by food category was conducted, as outlined in 3.1, ‘Product Categories and Freshness’. This included re-weighting of the Morrisons sample to reflect the increased proportion of Morrisons customers located in the ‘North West’ and ‘Yorkshire and the Humber’ regions.

To establish how Morrisons users populate their online grocery baskets (with respect to basket stability, price savviness and product discovery), an analysis of the page from which products were added to basket was performed, as outlined in Sect. 3.1, ‘Basket Stability and Price Sensitivity’.

3.1 Sample Preparation and Processing

Product Categories and Freshness. The Morrisons Sample consisted of 986,973 transacted food and drink items from 41,201 users/households obtained using the Google Query Explorer API. The sample was selected to broadly mimic the methodology used by the ONS [5] by partitioning dates into weeks and then taking randomly selected dates corresponding to each of the days of the week to form a sample "week" for comparison with the Living Costs and Food Survey April 2015 to March 2016 (LCF 2016) (Table 1).

The LCF 2016 "population proxy" collected by the ONS consisted of 4,760 households who were asked to document their spending on food over a two week period. The aggregated results were then reported as spend in £p.w. and divided by product category, broadly inline with the Classification of Individual Consumption According to Purpose (COICOP) codes [44]. The socio-demographic profiles of respondents were collected to map the results of this sample back to the parent population. The LCF forms the most complete estimate of national spending on food and drink, despite the relatively small sample size. For our study, a subsample of the COICOP codes was used to group foods into larger categories, to reduce the number of poorly categorised foods (e.g. ready meals containing a number of ingredients). The selected categories corresponding to major food groups are detailed in Table 2.

31,721 Morrisons products corresponding to those transacted by consumers in the sample time period were labelled with one of the categories in Table 2 and were also labelled 'fresh' or 'not fresh'. The revenue (pence per household p.w.) was then calculated for each product category. Table 3 shows that there is a clear discrepancy between the distribution of Morrisons customers and the national averages. Re-weighting of the sample was thus performed using the ONS mid-year population estimates [36] to re-calibrate the Morrisons sample. Due to the categorical non-ordinal nature of the data, comparison with the online and offline LCF 2016 spending data was done using the 2-sample χ^2-test of homogeneity.

Basket Stability and Price Sensitivity. A random sample of 195 million products transacted by Morrisons customers between 1 February 2016 and

Table 1. Morrisons sample composed in line with the technique used by LCF 2016

Day of week	Selected date
Monday	23/11/2015
Tuesday	28/04/2015
Wednesday	11/11/2015
Thursday	02/07/2015
Friday	30/10/2015
Saturday	29/08/2015
Saturday	07/02/2016

Table 2. Composite food categories derived from COICOP codes

Category label	Products included
Bread & cereals	Bread, pasta, lentils, pulses, savoury biscuits and buns, breakfast cereals, corn-based products, quiches, pastries, non-meat pies
Fruit & veg	Fresh, dried, frozen and processed fruit and veg, including potatoes
Meat	Fresh, cured and frozen meat products including meat pies
Fish	Fresh, processed and frozen fish including battered fish
Dairy & eggs	Fresh and dried milk, cream, yoghurts, eggs and dairy substitutes
Confectionary	Biscuits, cakes, sweet buns, sweets, chocolate, ice-cream, jelly, sugar, jams and sugared spreads
Non-alc. drinks	Beverages not containing alcohol, including tea and coffee
Other	Including soup, seasoning, baby food, butter, vegetable oils and protein food replacements

Table 3. Proportion of Morrisons' customers by location compared to 2015 mid-year population estimates [30,36].

Region	Morrisons sample	Mid-year 2015	Diff
West Midlands	13.9%	8.7%	+5.2%
South East	4.7%	15.3%	−10.6%
North West	13.9%	10.9%	+3.0%
East	3.0%	10.2%	−7.2%
Yorkshire and The Humber	21.5%	7.8%	+13.7%
South West	5.7%	9.2%	−3.5%
East Midlands	11.3%	7.5%	+3.8%
London	25.4%	13.4%	+12.0%
North East	0.3%	3.6%	−3.3%
Scotland	0.1%	8.8%	−8.7%
Wales	0.2%	4.6%	−4.4%

1 February 2017 was obtained using the Google Query API from Morrisons' Google Analytics account. The 'ga:productListName', 'ga:pagePath' and 'ga:eventLabel' dimensions were used to characterise products in terms of basket stability and price sensitivity. A summary of the factors used to characterise each group are shown in Table 4. Aggregated revenue was used to calculate average basket compositions by stability and price sensitivity.

Table 4. Factors characterising price sensitive, stable and disrupted adds to basket

Price sensitive	Stable	Disrupted
Offers	Shopping list	¬Shopping list
Flash sales	Favourites	¬Favourites
Sort by price ascending	Suggested order	¬Suggested order
	Previous order	¬Previous order

4 Results

4.1 Average Basket Value

The average basket value for the Morrisons sample per household per week (p.h.p.w.) was £33.56, considerably higher than the £20.93, as reported by the LCF 2016[1].

4.2 Product Categories and Freshness

Online LCF 2016 vs. Offline LCF 2016. The latest release of the LCF separated online grocery shopping by food category from offline spending. In order to assess whether there is a statistically significant difference in the distribution of products bought online and offline, we perform a 2-sample χ^2-test. The null hypothesis is defined as H_0^1 below.

Null Hypothesis 1 ($\boldsymbol{H_0^1}$). *The distribution of revenue between food categories for the LCF online and offline surveys 2016 are the same (Table 5).*

Table 5. LCF online vs. LCF offline - distribution of revenue by food category

df	7
α	0.005
$\chi^2_{0.005,7}$	20.278
χ^2	$286.968 \gg \chi^2_{0.005,7}$
ϕ	0.04
Result	Strong evidence to reject H_0^1

The results of the 2-sample χ^2 test comparing the distribution of revenue by food category in £0,000s indicate strong evidence to reject H_0^1 with $\chi^2 = 286.968 \gg \chi^2_{0.005,7} = 20.278$ (note, significant at 5%, 1% levels also). Due to the large sample sizes (revenues in the £m), χ^2 test results can be misleading.

[1] Calculated by dividing the total revenue for online grocery shopping by the number of households reported to have done online grocery shopping in the last 12 months.

For example, a £5 variation in revenue between food categories would seem significant compared with a total revenue of £50, but significantly less if the total revenue was £5ṁ. The samples have been scaled to help account for this, but we also present the sample-size ambivalent measure of “effect size”, given by:

$$\phi = \sqrt{\frac{\chi^2}{n}}$$

where n = total number of observations (total revenue) and standardised residuals given by:

$$standardised\ residuals = \frac{(observed - expected)}{\sqrt{expected}}$$

Cohen suggests that $\phi = 0.1$ indicates a small effect size, $\phi = 0.3$ a moderate effect size and $\phi = 0.5$ a large effect size [9]. As such, we see that despite the large χ^2, there is still only moderate evidence that the online LCF sample is significantly different from the offline LCF sample.

Examination of the standardised residuals reveals that the major contributors to the difference between the samples is the proportions of revenue from ‘Confectionary’ and ‘Meat’ (which were overweight in the offline sample) and the ‘Other’ and ‘Non-alcoholic drinks’ categories (which were overweight in the online sample) (Table 6).

Table 6. Morrisons sample vs. LCF online 2016 - distribution of revenue by food category

df	7
α	0.005
$\chi^2_{0.005,7}$	20.278
χ^2	$230.939 > \chi^2_{0.005,7}$
ϕ	0.02
Result	Evidence to reject H_0^2

Online LCF 2016 vs. Morrisons Online Sample. To assess whether the Morrisons online sample is representative of online grocery consumption at the national level (i.e. drawn from the ‘population at large’, we perform a 2-sample χ^2 independence test with null hypothesis:

Null Hypothesis 2 (H_0^2). *The distribution of revenue between food categories for the Morrisons group is the same as the LCF online survey 2016 (representing the national average).*

The results of the 2-sample χ^2 test (Table 7) comparing the distribution of revenue by food category in £’00s[2] indicate some evidence to reject H_0^2.

[2] £’00s was chosen since an increase in 1 frequency point corresponds to a reasonable c.2% swing in value for the category with least revenue, ‘Fish & Seafood’.

The effect size given by $\phi = 0.14$ indicates slightly more disparity between the two online samples than between the online and offline LCF samples. Examination of the standardised residuals reveals that the major contributors to the difference between the samples is the proportions of revenue from 'Bread and Cereals' (which is overweight in the Morrisons sample).

Table 7. Morrisons sample (weighted by region) vs. LCF online 2016 - distribution of revenue by food category

df	7
α	0.005
$\chi^2_{0.005,7}$	20.278
χ^2	$19.195 < \chi^2_{0.005,7}$
ϕ	0.14
Result	Insufficient evidence to reject H_0^3

Online LCF 2016 vs. Morrisons Online Sample (Re-weighted). As shown in Table 3, Sect. 3.1, there is significant evidence of bias in the location of the shoppers in the Morrisons sample. The next section considers the effect of re-weighting the sample by region. We therefore propose re-weighting by region, where the weighting factors are given by:

$$\frac{\text{National expected proportion for region}}{\text{Morrisons observed proportion for region}}$$

We then assess how effective this re-weighting of the sample by region is by proposing the hypothesis:

Null Hypothesis 3 (H_0^3). *The distribution of revenue between food categories is the same for the Region Re-Weighted Morrisons sample and LCF online survey 2016.*

Finally, noting that the ONS data detailing spend on groceries by region does not separate online and offline expenditure, we propose a further re-weighting using data for the proportion of the population of each region who have used the internet in the past 3 months (as at 2016) [35].

Null Hypothesis 4 (H_0^4). *The distribution of revenue between food categories is the same for the region re-Weighted, Internet users by region re-weighted Morrisons sample and LCF online survey 2016.*

The results of the χ^2 test shown in Table 8 indicating that with re-weighting by region and internet use by region, the Morrisons sample is not significantly different from the Online LCF 2016 survey sample. There is however relatively little improvement over the re-weighting by region only. In summary, the unweighted, re-weighted by region and re-weighted by region and internet use by region are shown in Table 9.

Table 8. Morrisons sample (weighted by region and regional Internet use) vs. LCF online 2016 - distribution of revenue by Food category

df	7
α	0.005
$\chi^2_{0.005,7}$	20.278
χ^2	$19.093 < \chi^2_{0.005,7}$
ϕ	0.13
Result	Insufficient evidence to reject H_0^3

Table 9. Summary of χ^2 and ϕ (effect size) for unweighted and re-weighted Morrisons sample compared to LCF online survey 2016

	Effect size (ϕ)	χ^2 result
Unweighted	0.02	Reject H_0
Weighted by region	0.1292	Fail to reject H_0
Weighted by region & Internet use	0.1288	Fail to reject H_0

Freshness. It is often reported that online grocery shopping is unsuitable for purchasing fresh produce due to the delay in delivering goods to the consumer and the fact customers cannot pick their own perishable products [24,26]. We test this assertion for our weighted sample against the LCF offline survey 2016.

Null Hypothesis 5 (H_0^5). *The distribution of revenue between fresh and non-fresh produce is the same for the re-weighted Morrisons online sample and the LCF offline survey 2016 (augmented with FFS 2015).*

Table 10 shows that there is sufficient evidence to reject H_0^5 at the 0.5% level, although this is accompanied by a small effect size, $\phi = 0.09$. Inspection of the distribution between fresh and non-fresh indicates that the proportion of fresh products is larger in the Morrisons weighted sample than in the LCF offline sample suggesting that there may in fact be a higher proportion of fresh products sold online.

Table 10. Morrisons (weighted by region and regional Internet Use) vs. LCF offline - distribution of revenue by freshness

df	1
α	0.005
$\chi^2_{0.005,7}$	7.879
χ^2	$9.542 < \chi^2_{0.005,7}$
ϕ	0.09
Result	Sufficient evidence to reject H_0^5

4.3 Basket Stability and Price Sensitivity

For the average basket over 2016, 39% corresponded to products added to basket from "stable" activities, such as via user-specified "favourites", saved "shopping lists" and "previous orders". Figure 1a shows that product adds from "favourites" were by far the most significant stabilising component, accounting for 93% of the 39% stable basket.

Figure 1b shows that 60% of the disruptive product adds came from engagement with the site search, with only 8% adds originating from the hierarchical site navigation. Engagement with offers yielded the second largest proportion of the disruptive product adds at 19%.

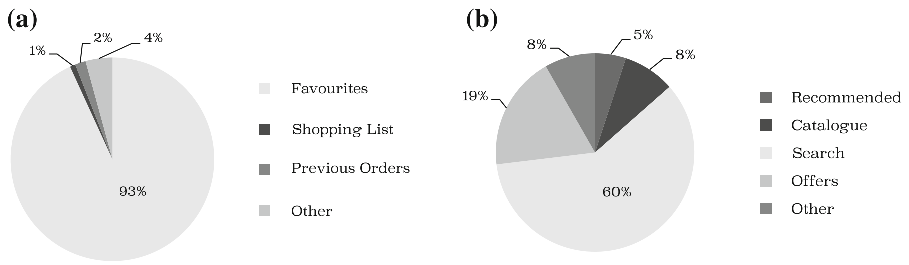

Fig. 1. (a) Breakdown of stable product adds (b) Breakdown of disrupted product adds

Figure 1b shows that the majority of product adds are considered "disrupted" since they required engagement with offers, featured/recommendations or site navigation to find relevant products.

"Price savvy" behaviours were defined as those emanating from engagement with offers, meal deals, flash sales and bundle deals, and those where search results were ordered by ascending price. Whilst engagement with offers constituted 23% of the >195 million "adds to basket" events over 2016, with just 1% were the result of customers actively sorting products by price (Fig. 2).

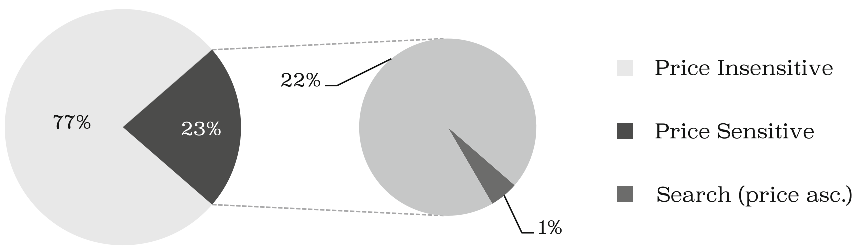

Fig. 2. Breakdown of price sensitive product adds

5 Discussion

The broad aims of this study were to assess the feasibility of using the Morrisons dataset for exploring national level online grocery behaviour; to assess whether there are differences in online and offline grocery consumption; and to examine in detail the features of online grocery that are thought to constrain or change behaviour. This was achieved by considering basket value; basket composition; basket stability and price sensitivity at the national level and in the Morrisons sample. The results and implications of the findings are discussed in Sects. 5.1–5.4.

5.1 Is the Morrisons Sample Representative of Online Grocery Shopping at the National Level?

When considering the composition of the average basket by product category, it was found that the Morrisons sample was significantly different from the LCF 2016 "national average" online group, but that re-weighting the Morrisons sample by region lessened the disparity such that it was no longer statistically significant at the 0.05% significance level (although could be rejected at the 1 and 5% levels). This was marginally improved by also weighting the sample by internet use by region to correct for the fact that there is currently no dataset that splits regional grocery consumption by online and offline channels.

The identification of a rich dataset which has potential to represent online grocery shopping behaviour at the national level presents exciting opportunities to expand understanding in a sparsely researched field. The extensive body of work to follow this study will not only contribute to the understanding of online grocery consumption, but will also serve as a scoping exercise to help guide research in the field.

5.2 Are There Differences in Consumption Behaviour Between Online and Offline Grocery Shopping in the UK?

Average Basket Value. The average basket value for the Morrisons sample was found to be £33.56 p.h.p.w., considerably higher than the national average of £20.93 p.h.p.w., as reported in the LCF 2016. This may indicate that Morrisons shoppers spend more than the national average, or that they shop less frequently, as reported by Robinson et al. [40]. Furthermore, the national estimate may be an underestimate of average basket size online, since it is based on the proportion of the population who have bought groceries online in the past 12 months, not those who do so on a regular basis.

Basket Composition. Comparison of the online and offline LCF 2016 samples indicated strong evidence of difference in the distribution of revenue among product types for online and offline grocery shopping, indicating that online consumption behaviour is different from offline. It was shown that there is some evidence that online consumers spend more on fresh products than offline, supporting the findings of Degeratu et al. [13], but contrary to popular wisdom that customers are reluctant to shop for fresh produce online [24,26].

5.3 Are Online Grocery Shoppers Price Sensitive?

Interrogation of Morrisons customer basket composition revealed that 23% of the average basket was populated using price sensitive behaviours, such as engagement with offers and sorting products by price. Whilst there are no comparative offline studies for Morrisons consumers, this finding indicates that the majority of products are added to basket without obvious attention to price and may indicate an overestimation by the retailer of consumer bargain seeking and price comparison behaviours. Furthermore, most "price savvy" adds to basket came from engagement with offers which are prevalent on the site and in search results. As such, the value-seeking behaviour appears to arise primarily from product placement on the part of the retailer. This may suggest an overestimation by the retailer of customer's value-seeking behaviour. This result is concordant with the findings of Urbany et al., who report that in the "offline-only era" retailers tended to 'overestimate the size of the consumer segment that regularly switches stores for price specials' [45]. Establishing how Morrisons' online price sensitive behaviour compares to in-store behaviour (and enable generalisation to the UK population) would of course require further empirical investigation.

5.4 Are Online Grocery Baskets Stable?

This study also provides initial findings regarding how consumers make product selections by identifying the page on-site from which products were added to basket. It was shown that 39% of product additions emanate from stable behaviours, such as using previous orders and saved favourites to populate shopping baskets. The vast majority of stable product adds are done so from saved favourites. The majority of all product adds to basket however result from "disrupted" activities, such as site search or engagement with retailer promotions. The high proportion of unstable product adds to basket could indicate that the retailer is overestimating the difficulty in up-selling and infiltrating online baskets.

6 Conclusions

This study established that there are differences in consumption behaviour between online and offline grocery shopping in the UK in terms of basket composition by product category. It also provides some evidence that online consumers spend more on fresh products than offline, contrary to popular wisdom that customers are reluctant to shop for fresh produce online [24,26]. This paper showed that with re-weighting to correct for regional bias, a rich dataset of Morrisons online consumer data has potential to represent national level behaviour. The dataset consists of individual level consumption behaviour and includes features such as location, time, device and the names of products purchased. This finding has huge implications for expanding the understanding of online grocery consumption at the national level. To date, the majority of quantitative empirical work has been conducted in non-naturalistic experimental settings and with small sample sizes. In contrast, the identification of this rich dataset facilitates analysis of millions of real-world transactions.

6.1 Limitations

The main limitations of this study resided in the lack of offline Morrisons data or high quality online and offline datasets to compare with to conclude whether there have been systematic shifts in consumer behaviour in the online grocery shopping era. Whilst the size of the sample used in this study was vastly bigger than that used by the ONS to estimate national level behaviour, it was still only a small proportion of the wealth of data collected by Morrisons and other retailers.

The assignment of product categories and freshness was not trivial, due to the inconsistent way CIOCOP and LCF surveys report on food categories. Some categories (e.g. fruit and vegetables) were well defined in terms of freshness, whilst others (such as meat and fish) were not divided into fresh and frozen/processed products. Furthermore, complex products containing multiple food groups were difficult to classify and no clear advice for doing so exists in the LCF e-commerce.

6.2 Future Work

Future work will identify periodicity in sales revenue; variation in behaviours by location and device use; and how the capacity to edit baskets affects basket composition. The discovery that a high proportion of product adds to basket emanate from unstable sources, such as site search also motivates qualitative investigation. This would involve working with shoppers to ascertain how they engage with online shopping, the reasons they give for their behaviours and further investigating sources of fundamental change in consumption practices. Employing insights from qualitative and quantitative studies will allow theoretical and predictive models of consumption behaviour to be developed.

References

1. Anesbury, Z., Nenycz-Thiel, M., Dawes, J., Kennedy, R.: How do shoppers behave online? An observational study of online grocery shopping. J. Consum. Behav. **15**, 261–270 (2016)
2. Bagozzi, R.: The legacy of the technology acceptance model and a proposal for a paradigm shift. J. Assoc. Inf. Syst. **8**(4), 244–254 (2007)
3. BI Intelligence: Mobile sales drive unexpected UK e-commerce growth. Business Insider UK (2017). http://uk.businessinsider.com/mobile-sales-drive-unexpected-uk-e-commerce-growth-2017-1

4. Briesch, R.A., Chintagunta, P.K., Fox, E.J.: How does assortment affect grocery store choice? J. Mark. Res. **46**(2), 176–189 (2009)
5. Bulman, J.: Detailed household expenditure by countries and regions, UK: Table A35. Office for National Statistics (2017). https://www.ons.gov.uk/peoplepopulationandcommunity/personalandhouseholdfinances/expenditure/datasets/detailedhouseholdexpenditurebycountriesandregionsuktablea35
6. Business: Ocado profits rise after 'robust trading'. BBC News (2017). http://www.bbc.co.uk/news/business-38806573
7. Chu, J., Chintagunta, P.K., Cebollada, J.: A comparison of within-household price sensitivity across online and offline channels. Mark. Sci. **27**(2), 283–299 (2008)
8. Cobb, C., Hoyer, W.: Direct observation of search behaviour in the purchase of two nondurable products. Psychol. Mark. **2**, 161–179 (1985)
9. Cohen, J.: Statistical Power Analysis for the Behavioral Sciences, 1st edn. Taylor and Francis, Milton Park (1988)
10. Crewe, L., Lowe, M.: Gap on the map? Towards a geography of consumption and identity. Environ. Plann. A **27**(12), 1877–1898 (1995)
11. Danaher, P.J., Wilson, I.W., Davis, R.A.: A comparison of online and offline consumer brand loyalty. Mark. Sci. **22**(4), 461–476 (2003)
12. Dawes, J., Nenycz-Thiel, M.: Comparing retailer purchase patterns and brand metrics for in-store and online grocery purchasing. J. Mark. Manag. **30**(3–4), 364–382 (2014)
13. Degeratu, A.M., Rangaswamy, A., Wu, J.: Consumer choice behavior in online and traditional supermarkets: the effects of brand name, price, and other search attributes. Int. J. Res. Mark. **17**(1), 55–78 (2000)
14. Dickson, P., Sawyer, A.: The price knowledge and search of supermarket shoppers. J. Mark. **54**, 42–53 (1990)
15. Edwards, J.: Retail sales in great britain - Office for National Statistics (2017). https://www.ons.gov.uk/businessindustryandtrade/retailindustry/bulletins/retailsales/mar2017\#whats-the-story-in-online-sales
16. Morrisons Employee: Interview with Morrisons (technical) employee (2016)
17. Evanschitzky, H., Iyer, G., Hesse, J., Ahlert, D.: E-satisfaction: a re-examination. J. Retail. **80**(3), 239–247 (2004)
18. Hand, C., Dall'Olmo Riley, F., Harris, P., Singh, J., Rettie, R.: Online grocery shopping: the influence of situational factors. Eur. J. Mark. **43**(9/10), 1205–1219 (2009)
19. Hansen, T.: Consumer adoption of online grocery buying: a discriminant analysis. Int. J. Retail Distrib. Manag. **33**(2), 101–121 (2005)
20. Huang, Y., Oppewal, H.: Why consumers hesitate to shop online. Int. J. Retail Distrib. Manag. **34**(4/5), 334–353 (2006)
21. IGD.com: Larger stores rally as UK food and grocery forecast to grow by 10% over next five years (2016). http://www.igd.com/About-us/Media/IGD-news-and-press-releases/Larger-stores-rally-as-UK-food-and-grocery-forecast-to-grow-by-10-over-next-five-years/
22. Julka, H.: Tech in Asia - connecting Asia's startup ecosystem (2016). https://www.techinasia.com/11-reasons-online-grocery-startups-failing-indias-small-towns
23. Keifer, S.: Why the internet is forcing retailers to rethink their supply chains. Computer Weekly, p. 15 (2013). http://www.computerweekly.com/opinion/Why-the-internet-is-forcing-retailers-to-rethink-their-supply-chains

24. Kestenbaum, R.: Why online grocers are so unsuccessful and what Amazon is doing about it. Forbes.com Retail (2017). https://www.forbes.com/sites/richardkestenbaum/2017/01/16/why-online-grocers-are-so-unsuccessful-and-what-amazon-is-doing-about-it/\#6718c6f77f56
25. Maity, M., Dass, M.: Consumer decision-making across modern and traditional channels: e-commerce, m-commerce, in-store. Decis. Support Syst. **61**, 34–46 (2014)
26. Marino, J.: The online grocery business faces a unique challenge. Business Insider UK (2015). http://uk.businessinsider.com/grocery-delivery-challenges-2015-3?r=US&IR=T
27. Marr, B.: Big data and shopping: how analytics is changing retail. LinkedIn (2015). https://www.linkedin.com/pulse/big-data-shopping-how-analytics-changing-retail-bernard-marr
28. Michaud Trevinal, A., Stenger, T.: Toward a conceptualization of the online shopping experience. J. Retail. Consum. Serv. **21**(3), 314–326 (2014)
29. Miller, D.: A Theory of Shopping, 1st edn. Polity Press, Cambridge (1998)
30. Morrisons.com: Groceries.morrisons.com (2017). https://groceries.morrisons.com/webshop/startWebshop.do
31. Murphy, A.J.: (Re)solving space and time: fulfilment issues in online grocery retailing. Environ. Plann. A **35**(7), 1173–1200 (2003)
32. Neville, S.: Tesco: who says it's hard to make home delivery profits? Independent (2014). http://www.independent.co.uk/news/business/news/tesco-who-says-its-hard-to-make-home-delivery-profits-9153185.html
33. Newman, D.: Reimagining the holiday experience with big data. IBM THINK Marketing (2016). https://www.ibm.com/think/marketing/reimagining-the-holiday-experience-with-big-data/
34. Office for National Statistics: Retail sales (2012). http://webarchive.nationalarchives.gov.uk/20160105160709/, http://www.ons.gov.uk/ons/dcp171778_263543.pdf
35. Office for National Statistics: Internet users in the UK: 2016 (2016). https://www.ons.gov.uk/businessindustryandtrade/itandinternetindustry/bulletins/internetusers/2016
36. Office for National Statistics: Mid-2015 (2017). https://www.ons.gov.uk/peoplepopulationandcommunity/populationandmigration/populationestimates/datasets/populationestimatesforukenglandandwalesscotlandandnorthernireland
37. Rafiq, M., Fulford, H.: Loyalty transfer from offline to online stores in the UK grocery industry. Int. J. Retail Distrib. Manag. **33**(6), 444–460 (2005)
38. Reichheld, F.F., Sasser, W.E.: Zero defections: quality comes to services. Harv. Bus. Rev. **68**(5), 105–111 (2000)
39. Ritzer, G.: The McDonaldization of Society 6, 1st edn. Sage, Thousand Oaks (2011)
40. Robinson, H., Dall'Olmo Riley, F., Rettie, R., Rolls-Willson, G.: The role of situational variables in online grocery shopping in the UK. Mark. Rev. **7**(1), 89–106 (2007)
41. Rohm, A.J., Swaminathan, V.: A typology of online shoppers based on shopping motivations. J. Bus. Res. **57**(7), 748–757 (2004)
42. Szymanski, D.M., Hise, R.T.: E-satisfaction: an initial examination. J. Retail. **76**(3), 309–322 (2000)
43. Thachenkary, C., Chatterjee, S., Katz, J.: Successful product characteristics for electronic commerce: a taxonomy of transaction types. In: 1997 Fourth International Workshop on Community Networking Processing (1997)
44. United Nations Statistics Division: United Nations statistics division - classifications registry (2017). https://unstats.un.org/unsd/cr/registry/regcst.asp?Cl=5

45. Urbany, J.E., Dickson, P.R., Sawyer, A.G.: Insights into cross- and within-store price search: retailer estimates vs. consumer self-reports. J. Retail. **76**(2), 243–258 (2000)
46. Wajcman, J.: Pressed for Time: The Acceleration of Life in Digital Capitalism, 1st edn. University of Chicago Press, Chicago (2015)
47. Wang, R.J.H., Malthouse, E.C., Krishnamurthi, L.: On the go: how mobile shopping affects customer purchase behavior. J. Retail. **91**(2), 217–234 (2015)
48. Worldpanel, K.: Global e-commerce grocery market has grown 15% to 48bn. Kantar Worldpanel News (2016). https://www.kantarworldpanel.com/global/News/Global-e-commerce-grocery-market-has-grown-15-to-48bn

Open Data as a New Commons. Empowering Citizens to Make Meaningful Use of a New Resource

Nicola Morelli[1], Ingrid Mulder[2], Grazia Concilio[3], Janice S. Pedersen[4], Tomasz Jaskiewicz[2], Amalia de Götzen[1](✉), and Marc Arguillar[5]

[1] Aalborg University, Copenhagen, Denmark
{nmor,ago}@create.aau.dk
[2] Technische Universiteit, Delft, Netherlands
[3] Politecnico di Milano, Milan, Italy
[4] Antropologerne ApS, Copenhagen, Denmark
[5] I2CAT, Barcelona, Spain

Abstract. An increasing computing capability is raising the opportunities to use a large amount of publicly available data for creating new applications and a new generation of public services.

But while it is easy to find some early examples of services concerning control systems (e.g. traffic, meteo, telecommunication) and commercial applications (e.g. profiling systems), few examples are instead available about the use of data as a new resource for empowering citizens, i.e. supporting citizens' decisions about everyday life, political choices, organization of their movements, information about social, cultural and environmental opportunities around them and government choices. Developing spaces for enabling citizens to harness the opportunities coming from the use of this new resource, offers thus a substantial promise of social innovation.

This means that open data is virtually a new resource that could become a new commons with the engagement of interested and active communities. The condition for open data becoming a new commons is that citizens become aware of the potential of this resource, that they use it for creating new services and that new practices and infrastructures are defined, that would support the use of such resource.

Keywords: Open data · Social innovation · Commons

1 Introduction

The term commons was used in medieval England to indicate parts of land which were openly accessible to everybody in a community. By extension the term was lately used to indicate any resource (land, forests, water basins, or pastures) that was freely available but controlled by a community, through a set of rules that regulate the access in order to create a sustainable use of the resource.

I. Kompatsiaris et al. (Eds.): INSCI 2017, LNCS 10673, pp. 212–221, 2017.
https://doi.org/10.1007/978-3-319-70284-1_17

The open availability of such resources is of course not a sufficient condition for them to survive as a commons and be used in a sustainable way. Without a community taking care of a commons pasture, for instance, this land would be neglected or someone would claim its ownership. In order for a commons to persist and grow, it is necessary that a defined group of people develop their own distinctive social practices and bodies of knowledge for managing the resource (Ostrom, 1990).

This is also true for "modern" resources, or "new commons" (Hess, 2008), of which a very famous example is the widely used open software platform known as GNU/Linux (Bollier, 2014). If a commons can be defined as *a resource* + *a community* + *a set of social protocols* (Bollier, 2014, p. 15), what is still to be designed in order to define open data as a new commons? Why it is important do definit it as such? What are the challenges and the opportunities in the current socio-technical systems?

The latest development of ICT, Internet of things and communication technologies are making it possible to collect a very large amount of data. Computing capabilities are in turn making it possible to coordinate the collection of those data in coherent data sets that can provide information on many aspects of our life, our cities, our environment, social life and healthcare condition. This increases our capability of monitoring our environment.

The potential of this phenomenon is still largely unexplored, because if, on one hand, the input side is becoming very diversified and providing rich data sets, on the other hand the output side is only used by few stakeholders, that have identified good opportunities for their business. Many businesses and institutions take advantage of large and typically closed data sets. For example, Google's core business is to collect data about users of its services such that it can better target its advertisements to individuals; insurance companies collect data about their customers to assess insurance risks; and tax authorities use big data to identify cases of tax fraud. On the other hand, there are many data sets that do not contain sensitive information and that already are, or could be, made open by governments and companies for public use. However, non-experts do not have the knowledge and skills needed to take advantage of such large data sets in the way that large institutions or companies do.

1.1 The Input Side

An increasing presence of microprocessors, sensors and similar devices to control several aspects of urban and business life, such as traffic, access to public buildings or environmental conditions is generating a large amount of data that can be easily collected and made available. Without touching the delicate issue of privacy, related to this data collection, which opens ethical questions that cannot be ignored, the availability of those data could also be and in some case is being used as an opportunity to control crucial aspects of our environments, from parking to traffic regulation, environmental monitoring and garbage collection.

The miniaturization of all this technology made it possible to have wearable devices that can also be used to consciously track ourselves: we can monitor our body in many different ways, record how many steps we have taken in one day and so on. This is what is called the quantified self, a trend in continuous expansion. Such trends deal with very private information, but we often decide to share those data among friends or even publicly, adding information that also others can use (e.g. on weather conditions, condition of the area where we have just been running, etc.).

Another important source of data comes from social networks, in which users are more or less voluntarily providing data about their preferences, geographical positioning, political orientation, and even private life. Those data have so far been used for commercial purposes, because they give chances for accurate profiling of social network users and therefore make it possible to generate highly personalized commercial offers. However social networks have also been used for different purposes; they have been working as an aggregation tool for movements, local initiatives, or political groups. Social network users often provide on-time qualitative data about places, weather, transport problems, healthcare crisis, emergency, security issues, catastrophes, etc.

The integration of those two kinds of datasets, the automated data input coming from micro-controllers disseminated in our daily environments and the voluntary data input provided by users could improve our understanding of the context in which we live in, because it would link quantitative data, often hard to interpret, to qualitative data, which are sometimes imprecise, but very useful to tell a narrative, allowing us to better navigate in such a complex context, giving meaning to numbers. This paper argues that such integration presents opportunities to consider data as a new public resource. For this resource to be considered as a commons a new awareness should be raised on the output side, to make sure that more people – citizens, organisations, interest groups, public administrations – be interested in using this resource and developing tools and shared practices to take care of this resource

1.2 The Output Side

The initiatives to make datasets available to the broader public are becoming more and more frequent. Several cities, such as Copenhagen, Rotterdam, Glasgow, Barcelona etc. have opened their databases, several datasets are now freely accessible on the Web and can be used by anyone.

The question is how many stakeholders have the capability to use those data for relevant applications. The use of those data in fact requires:

1. technical skills, to create applications based on the available data,
2. creative capabilities, to understand which needs those datasets can address and for which part of our society,
3. data literacy, to pose answerable research questions and interrogate the available data (or create a dataset).

It has already been mentioned that some early examples of the use of the datasets concern control systems (e.g. traffic, meteo, telecommunication) and commercial applications (e.g. profiling systems); few examples are available about the use of such data as a new resource for empowering citizens, i.e. supporting citizens' decisions about everyday life, political choices, organization of their movements, information about social, cultural and environmental opportunities around them, government choices. Developing means for enabling citizens to harness the opportunities coming from the use of this new resource, offers thus a substantial promise of social innovation.

2 Context: The Open4Citizens Project

Looking back at open data, although this resource has been seized by many companies to be used for commercial purposes, its potential is still widely unexplored and could still have large scope for applications and uses that can directly benefit communities, even outside the market mechanisms. Open data can be the base of a new generation of public services that is directly defined, designed and used by citizens. Open data can become a commons if communities are built around them, that create sets of practices and rules to use this resource in the most sustainable way. This is the aim of the EU-Funded Open4Citizens project (www.open4citizens.eu), which is experimenting on new forms of collaboration between citizens, public authorities, interest groups, local businesses and IT experts, with the aim of (a) generating new services, (b) aggregating communities around the resource (c) creating new infrastructures for the use of open data.

The project focuses on Open Data as a new shared resources and aims at generating the conditions for this resources to be used and managed as a commons. More specifically the project is developing strategies and tool with the aim of:

- Aggregating a community that takes care and uses the resource
- Developing a set of practices for using the resource
- Infrastructuring the community with a set of tools that facilitate the access to the resource and its usage

The project is articulated in 5 pilots in Barcelona, Copenhagen, Karlstad, Milano and Rotterdam. Each of these pilots works on different challenges (Table 1) to define different solutions, within a shared framework of design processes, tools and methods to increase citizens' awareness of open data and engage them in the creation of new solutions for their everyday problems.

Table 1. Competences and skills in the ecosystem of stakeholders participating in the hackathon

Stakeholder	Knowledge/experience/skills
Citizens	Personal/daily life experience
Interest groups	Domain experts, issue experts
IT experts/hackers	Technical knowledge
Public authorities	Institutional knowledge (regulations, public policies)
Data experts	Data literacy (scraping, cleaning, visualizing)
Entrepreneurs	Business knowledge
Data owners	Data format, potential of existing data, available data

This paper focuses on the creation of the needed infrastructure for the exploitation of open data as a new commons, and on the aggregation of communities (ecosystems) of users highlighting the hackathon process as an instrumental mean to facilitate the creation of this culture.

2.1 Aggregating Communities

Citizens' participation in the Open4Citizens pilots and their awareness in the use of open data are supported through the organization of *hackathons*.

Hackathons are a common practice among information technology (IT) enthusiasts and there is an extensive literature reporting on hackathon experiences (Linnel, 2014; Tanenbaum, 2014; Hecht, 2014). They consist of a 2–3 day *pressure cooker* event in which groups of IT experts collaborate to develop new applications. They are a playground for IT experts, that can propose new solutions, usually very innovative from a technological point of view. However these solutions rarely address citizens' needs, because of the nature of the hackathon, which is gathering experts coders, neglecting though the prospective users of such solutions/apps (i.e. those who are able to propose meaningful problems).

The purpose of the Open4Citizens hackathons is to extend the participation to non-IT experts, that means common citizens, interest groups, public servants and business organisations, which can together work on a commonly defined problem area. For this reason the phase of organization of each hackathon is critical, to ensure the participation of people with different skills, motivations and knowledge. Each Open4Citizens hackathon aspires to put together a local community that shares values, interests and motivations in specific problematic areas, thus creating an *ecosystem* to generate or consolidate the demand for open data or to organize crowdsourced gathering of new open data.

In order for a common good to benefit members of a community, there is a need for a shared understanding of the value of this resource as well as clear guidelines for its use. The Open4Citizens hackathon approach focuses on placing ordinary but engaged citizens at the centre of the development of open-data driven solutions to societal challenges that are important to them. This approach ensures that collaboration between hackathon participants with different skill sets and interests is appropriately facilitated such that there is a shared understanding of the potential of open data, all participants have the opportunity to manipulate open data and the contributions of all participants to the co-creation of solutions are valued. Some of the main competences and skills for the hackathons in Open4Citizens are listed in Table 2.

Table 2. Challenges in each Open4Citizens pilot

Pilot	Thematic areas
Copenhagen (DK)	Migration, integration
Karlstad (SE)	Healthcare, quantify self
Rotterdam (NL)	Parks and common spaces
Barcelona (SP)	Healthcare, culture
Milano (IT)	Transparency in public decisions for urban transformation

2.2 Building Practices: The Hackathon Process

The hackathon approach is meant to generate a set of practices to support citizens participation. Unlike the better known hackathons, the O4C hackathons aggregate quite heterogeneous knowledges and cultures; therefore the process needs to be accurately designed to reconcile dissimilar perspectives. The O4C team organized the hackathon as a process including:

- An initial **pre-hack** phase, characterized by an intense activity of dataset identification, verification and specification, running in parallel to the definition of the challenge(s) that the hackathon is supposed to tackle.
- The **hackathon event,** a *pressure-cooker* co-design event involving all the relevant stakeholders; and
- A **post-hack** phase, including follow-up activities to implement technical and business aspects of the outcomes of the post-hack.

2.3 Infrastructuring for Using Open Data

The Hackathon Starter Kit

In order to support participation in a hackathon by anyone with an interest in open data or in the societal challenges to be solved, a starter kit is proposed, that includes tools (Fig. 1) that support citizens in the various steps in a hackathon process while guiding them progressively from inspired idea to concrete solutions.

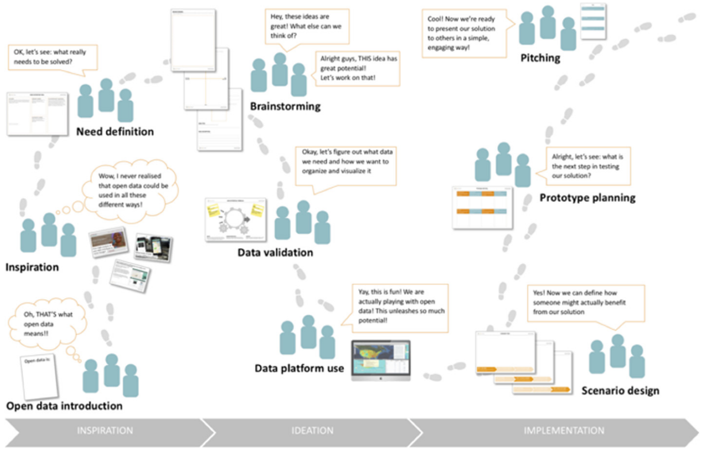

Fig. 1. Hackathon starter kit

Each hackathon process involves an interplay between three types of activities, which are: (1) “inspiration”, understanding the opportunities on the intersection of citizens’ needs and open data current or potential availability, (2) “ideation”, i.e., creatively addressing these opportunities, and (3) “implementation”, translating the ideas into concrete app or service concepts and prototypes. Physical (paper) tools complement and interact with the Open4Citizens platform (www.opendatalab.eu), that also includes visualization development tools.

The Open4Citizens (O4C) Platform

The O4C platform is a digital service for sharing, aggregating, curating and supporting open data and data-related knowledge, know-how and best practices within and across engaged communities. The platform serves as a backbone underlying the hackathons in the five pilot cities involved at this stage, supporting a shared understanding across different societal contexts and thematic challenges to be addressed using open data. It is intended to have low skill entry point and be user friendly so that citizens, hackathon organisers, facilitators and other stakeholders are able to use it, even if they have no previous experience of working with open data. It allows for shared open data usage across locations and for feedback to be gathered regarding its usability.

The O4C platform is being adapted in line with learning from hackathons to support the development of a network of pilot OpenDataLabs in the five project locations: Barcelona, Copenhagen, Karlstad, Milan and Rotterdam. Pilot hackathons and related activities provide test cases for the development of a shared vision of OpenDataLabs including the role that the O4C platform should play in underpinning the lab network.

The OpenDataLab

OpenDataLab represents the main conceptual and operational output of the Open4Citizens project. In the ongoing project an initial OpenDataLab has been conceptualized, drawing inspiration from FabLabs, where citizens' good ideas about new products can find knowledge and technologies to support the realisation of their ideas. Likewise, in OpenDataLabs citizens, groups or associations will be able to develop their ideas for new services using existing open data. The Open4Citizens project has already produced some deeper understanding of what an OpenDataLab can be. Surely it represents the space where an ecosystem of citizens, ICT experts and Public service managers is activated to generate solutions to urban problems by using open data. Furthermore the OpenDataLabs will be supported by the online platform presented in the previous section.

3 The First Hackathons' Cycle

As already mentioned, the O4C project is organized in 5 pilots where the above mentioned tools and methods developed by the consortium are tested and evaluated in two hackathon cycles. The evaluation of the first cycle has been the building blok for the design of the second one, that will run in fall 2017.

The experiments on open data in each pilot had to be articulated along three main dimensions: (1) the availability of data, (2) the nature of the societal challenge to be addressed and (3) the creation of a community, i.e. people with different interests, motivations, skills and knowledge that could experiment on practices to work with open data.

Those dimensions have been important to define the thematic areas for the activities in each pilot, which are synthesized in Table 2.

If we focus on the outcome of the hackathons, different results can be highlighted on the basis of the variation of the three above mentioned dimensions: data, challenges and people. The proposed solutions ranges from apps, public services, visualisations, online plaforms, physical places, online portals etc. and therefore also their level of development varies from concepts to prototypes, that are being developed as new services or as part of existing services.

As examples, *Seek a seeker* is a platform to connect potential hosts and asylum seekers. Locals (Danes and former asylum seekers) “adopt” or host a new asylum seeker for a period of time. The platform uses open data in a crowdsourced way. *Cantieri miei* is a service that uses existing open data to help citizens affected by the metro construction works in Milano in getting compensation payments. *PolenCat* is an app (under development) which integrates an open database of historical and current pollen levels with the user-contributed medical and geo-localisation data, and recommends personal daily routes that avoid or mitigate those areas expected to have high concentrations of pollens and other allergens to which the user is susceptible (see Fig. 2).

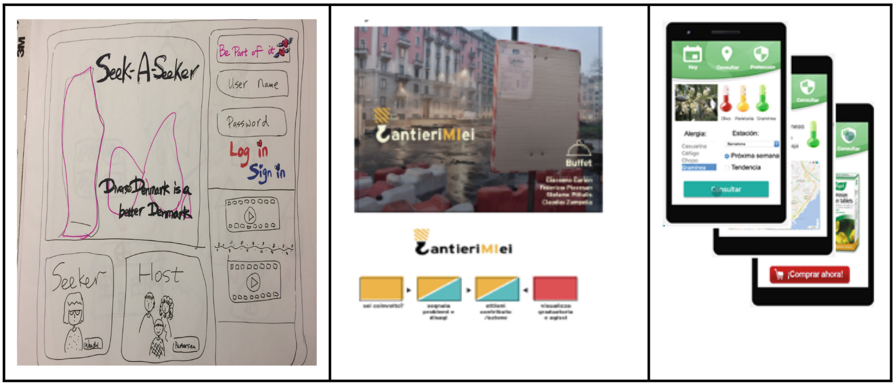

Fig. 2. Three solutions: seek a seeker - the concept of a platform from Copenhagen; CantieriMiei - a service under develpement from Milano; PolenCat, an app under development from Barcelona.

In these three examples the use of open data has also varied a lot: from imagining crowd sourced data sets in the Copenhagen pilot to a specific data set that was open from the municipality of Milan explicitely for the hackathon, to the use of available open datasets in the Barcelona’s pilot. The question of transparency in Milan, in fact, was directly referring to the availability of data to describe specific problems within the city. In that case the main effort has been focused on gathering a number of relevant stakeholders, which own the data and could suggest meaningful uses of them. In other pilots, where the problem was more widely defined (e.g. immigration/integration in Copenhagen) the proposed challenge did not directly point of relevant datasets, and the work of the participants has consisted in defining possible solutions on the basis of datasets to be crowdsourced.

All in all it is clear that significant work is needed in the pre-hack phase and into the hackathon itself to make data available in relevant formats and to provide guidance on how to use it. In this regard, the OpenDataLabs could have a crucial role and the first hackathon cycle has been a test-bed for developing different possible OpenDataLabs and identifying how they can become a sustainable, linked network.

4 Discussion

It could be argued that the current state of play regarding useful access to and use of open data by anyone wishing to do so is more reminiscent of the old 'tragedy of the commons' (Hardin 2009) idea than we would like. This can more appropriately be defined as the 'tragedy of open access regimes' (Daly and Farley 2011) when relating to empowering citizens to make appropriate use of open data. In this understanding, a resource that should be available to anyone can in fact only be used by a small subset of individuals or organisations with the skills and capacity to benefit from its potentials. In the case of open data, the Open4Citizens project suggests that the skills needed are technical skills, e.g. for creating apps or other means of open data integration in solutions, and creative capabilities, e.g. capabilities to identify opportunities where open data could be useful and imagining how they might be used.

The project team is testing how best to make these skills available to citizens or to support them in gaining them. By developing fit-for-purpose tools, guidelines for how to use them, and a support infrastructure for creating potentially impactful solutions to real-world challenges in urban service delivery, Open4Citizens seek to empower citizens to use the open data commons.

It has been demonstrated that understanding local successes in sustainable and equitable use of common resources provides lessons for use of large-scale commons (Ostrom et al. 1999). By learning from and comparing the facilitators and barriers relating to open data use for the common good in five different urban contexts across Europe, the project is laying the basis for an understanding of how an open data movement might be achieved. This movement is defined by the use of open data by any individual or community wishing to do so in a way that potentially benefits all citizens.

In order to learn from successes and challenges in open data use for the common good, Open4Citizens is working towards understanding the complex interactions at the local pilot level between (1) people interested in the potentials and practical applications of open data, (2) the nature of the open data landscape locally, i.e. how much data is open, how open it is in practice, whether the available open data is also useful and (3) the types of societal challenges that can appropriately be addressed using open data.

Beside the concrete solutions the Open4Citizens hackathons lead to, the project is also emphasising the potential of a hackathon event for fostering communities and open-data related knowledge, skills and know-how in these communities. Consequently, the OpenDataLab is a means to sustain and nurture these communities beyond hackathon events which are limited in duration.

5 Conclusions

The use of technology by human beings is always mediated (Miller & Slater 2000).

Technological elements such as open data, data manipulation tools and the fit-for-purpose OpenDataLab platform are only one, albeit essential, piece to establish an open data movement where open data truly become a common good. To this end, it is important to understand the needs of individuals and communities wanting to make use of this commons as well as the barriers and opportunities related to their open data use. It is also very important to lower the barriers to participation, so that contributing to this work isn't dependent on mastery of information technologies. For this reason, the knowledge base and guidelines being developed in this project regarding the interactions between analog tools and the digital, technological tools are key.

A central element in applying emerging lessons from local pilots to a larger scale European or even global open data movement involves understanding the nature of the commons: both the open data as a resource and the community or communities that want to use it. Linking these requires a shared understanding by the entire community (citizens, academics, the public sector and the private sector) of how open data can be used and to what end(s).

References

Ostrom, E.: Governing the Commons: The Evolution of Institutions for Collective Action. Cambridge University Press, Cambridge (1990)

Ostrom, E., Burger, J., Field, C.B., Norgaard, R.B., Policansky, D.: Revisiting the commons: local lessons, global challenges. Science **284**(5412), 278–282 (1999)

Hess, C.: Mapping the new commons. In: The Twelth Biennial Conference of the International Association for the Study of the Commons, Cheltenham, UK, 24–28 July 2008

Linnell, N., Figueira, S., Chintala, N., Falzarano, L., Ciancio, V.: Hack for the homeless: a humanitarian technology hackathon. In: 2014 IEEE Global Humanitarian Technology Conference (GHTC), pp. 577–584. IEEE, October 2014

Tanenbaum, K., Tanenbaum, J.G., Williams, A.M., Ratto, M., Resch, G., Gamba Bari, A.: Critical making hackathon: situated hacking, surveillance and big data proposal. In: CHI 2014 Extended Abstracts on Human Factors in Computing Systems, pp. 17–20. ACM, April 2014

Hecht, B.A., Werner, J., Raskar, R., Jouttenus, T.T., Khandbahale, S.S., Jouttenus, M., Bell, P.: The kumbhthon technical hackathon for nashik: a model for STEM education and social entrepreneurship. In: 2014 IEEE Integrated STEM Education Conference (ISEC), pp. 1–5. IEEE, March 2014

Daly, H.E., Farley, J.: Ecological Economics: Principles And Applications. Island press, Washington, D.C. (2011)

Hardin, G.: The tragedy of the commons. J. Nat. Res. Pol. Res. **1**(3), 243–253 (2009)

Miller, D., Slater, D.: The Internet: An Ethnographic Approach. Berg, Oxford (2000)

Involving Users in the Design of Sharing Economy Services

Asbjørn Følstad(✉), Dimitra Chasanidou, Ida Maria Haugstveit, and Ragnhild Halvorsrud

SINTEF, Oslo, Norway
asbjorn.folstad@sintef.no

Abstract. Involving users in the design of sharing economy services is important to realize the expected growth in this market. However, such involvement may be challenging due to the complexity and networked character of the service context. We present a case study showing how users' online feedback on novel design concepts may represent a viable approach to user involvement. In particular, the feedback provides insight into the strengths and weaknesses of proposed concepts as well as suggestions of relevance to the subsequent design process. On the basis of the case study, lessons learnt are discussed, as is needed future research.

Keywords: Sharing economy · User involvement · Online design feedback

1 Introduction

The sharing economy is expected to see substantial growth in Europe the coming years, with an estimated potential value of 570 billion Euros in 2025 [1]. The European Commission has pointed out the high potential for new businesses in this market [2].

To realize this growth, there is a need for sharing economy services that are seen as attractive and useful to a broad user group to ensure sufficient service uptake. For this purpose, a user-centred approach is essential throughout the stages of designing, implementing and evaluating sharing economy services [3].

However, involving users in such design and development is challenging. First, the resource demand for user involvement may be prohibitive for involving users sufficiently early in the process of ideation and concept formation. Second, novel sharing economy concepts may be difficult to represent so that they may be tried out directly by users.

In response to these challenges, we have explored the use of online design feedback for novel sharing economy concepts. The work is conducted as part of an innovation project on sharing economy services and business models, specifically addressing online redistribution markets. In this paper, we first discuss the challenges of involving users in the design of sharing economy services. Then, we present a case study of user involvement through an online system for design feedback. Finally, we present lessons learnt and suggest future challenges. The study contributes new knowledge on how an online approach to eliciting design feedback from users in response to visualizations of service concepts may benefit the design and development of sharing economy services.

I. Kompatsiaris et al. (Eds.): INSCI 2017, LNCS 10673, pp. 222–230, 2017.
https://doi.org/10.1007/978-3-319-70284-1_18

2 User-Centred Design of Sharing Economy Services

2.1 Sharing Economy Services

Sharing behaviour has existed as long as humankind. However, with the increased prevalence of Internet-based consumer devices, the sharing economy and collaborative consumption has emerged as a new way of accessing goods and services, where digital platforms are applied to connect the supply and demand side [4]. Botsman and Rogers [5] distinguish three main areas of collaborative consumption, including *collaborative lifestyles*, leveraging sharing of non-tangible assets such as skills and competencies, *product-service systems,* concerning the sharing of tangible assets such as consumer goods, and *redistribution markets*, addressing peer-to-peer distribution of used goods.

Sharing economy services typically concern the utilization of residual value, that is, making use of resources that would otherwise go to waste. Often, but not always, the end users of sharing economy services are non-professionals. With a digital platform connecting users, typically a provider/seller and a consumer/buyer, sharing economy services represent two-sided markets [6].

2.2 The Challenge of User Involvement in Sharing Economy Service Design

The current landscape of sharing economy services is characterized by a small number of large, transnational service providers such as AirBnB, Uber, some established providers dominating specific countries or regions, and a sprawling underwood of newcomer service providers aiming to capture markets shares from the established service providers and target new markets.

While future innovation in the sharing economy may depend on newcomer service providers, these are challenged in terms of involving users in the design process. As opposed to established service providers, access to users is challenging, and resources for involving users in design processes are scarce. Hence, in practical design and development projects user involvement may be conducted too late in the design process at a point where key design decisions have already been made. The risk of such late or inadequate user involvement may be that novel services fail to address key user needs, are not seen as adding value to users' everyday lives, and subsequently fail in the market [7].

2.3 The Challenge of Representing Sharing Economy Services

When involving users in the design of sharing economy services, representing these in an understandable manner to the users is critical. However, as sharing economy services typically are two-sided markets where independent user groups are connected through a digital platform, simulating service concepts is often not possible.

Service prototyping [8] is a feasible option in some service areas, such as in-store service provision. A service that implies geographical distance between customers, non-standard goods or services, non-professional users on both sided of the transaction, or networks of interacting services providers, is more challenging to prototype. In response, other – possibly more simplistic – means of representation are needed.

2.4 Visualizations and Idea Cards to Represent Service Concepts

Within the emerging discipline of service design, the challenge of representing service concepts have been given much attention. In particular, visualizations have become a much-used means of representing service concepts [9]. Often, rather complex service concepts may be represented through relatively simple visualizations in the form of images.

The use of brief description of ideas or concepts, often including an illustration or visualization, is another approach to represent larger sets of early-phase design concepts in order to process these further as part of the design process. Such idea cards have shown to be useful for gathering feedback from users [10].

2.5 Online Involvement of Users

Traditional user involvement practices involve face-to-face methods, such as workshops, interviews, focus groups, or evaluation sessions [11]. To reduce the resource demand of user involvement, while also allowing for the involvement of geographically dispersed users, researchers and practitioners have sought to leverage online environments for involving users in the different stages of the design process, through the use of general purpose environments such as Facebook groups [12] or dedicated environments for user involvement in design and innovation [13].

We have previously conducted online user feedback sessions for design processes towards novel interactive systems and services [14]. However, we have not previously tried this approach for services in two-sided markets such as the sharing economy.

3 Research Objective: Overcoming the Involvement Challenge

Building on existing research on the use of visualizations and idea cards to represent service concepts, as well as research on online user involvement in design processes, we in this study aim to explore how these means may be used to overcome the involvement challenge in the design of sharing economy services.

Hence, the research objective of this study is to gain experience and knowledge on how simple, visual representations of service concepts may be used for online elicitation of feedback from users when designing sharing economy services, and how such design feedback may support the subsequent design process. On the basis of these experiences, we should be better equipped to discuss how newcomer sharing economy service providers may establish user involvement in their design and development process.

4 Case Study: Involving Users in Design for Online Redistribution Markets

In response to the research objective, we conducted a case study [15] where a range of concepts for an online redistribution market were identified and presented for user feedback. In the following subsections, we present the case context, the identification and representation of design concepts, how users were engaged in evaluating the concepts, and the nature of the design feedback that was gained.

4.1 Scoping of the Case Study

The case study was carried out in January 2017 in collaboration with a provider of a relatively new online redistribution platform. This online market platform is designed through a "mobile first" approach, mainly oriented towards local-area transactions, and with the aim of providing a highly simplified process when selling and buying used items. From preliminary user insight studies, we had learnt that the cumbersome part of the service experience was the transaction process involving both the seller and the buyer, from their first contact until the item is sold. Key phases of this transaction process are (a) matching the prospective seller and buyer, (b) communication between the seller and buyer, (c) handover of goods and money, and (d) post-handover rating for social filtering. Challenges include easy filtering of irrelevant matches, e.g. due to geographical distance, inefficient or incomplete dialogues between sellers and buyers, transportation and handover challenges, and challenges motivating users to provide honest feedback following transactions. The case study involved the development of innovative concepts addressing these challenges, and gathering design feedback from the users.

4.2 Identifying and Representing the Design Concepts

A team of four researchers within the field of user-centred design individually drafted concepts targeting the identified challenges of the transaction process. The research team then discussed and negotiated the design concepts on the level of specific wording and illustrations for representing each concept. In addition, an external designer was invited to comment on the representation before finalizing the concepts.

The design concepts were elaborated in the form of digital idea cards, with an easy to understand illustration and a brief text motivating and outlining the concept. In total, a dozen design concepts were drafted; six of which were selected to be included in the study. An overview of the six concepts is provided in Table 1. An example design concept is presented in Fig. 1.

4.3 Involving Users in Online Design Feedback

To involve users in the design process, an online environment for design feedback (http://recordlivinglab.org) was used. The environment offers an easy way to gather user feedback from a broad set of users that are geographically dispersed. All participants were encouraged to provide feedback on each design concept in the form of

Table 1. Overview of design concepts

Transaction phase	Design concept
Matching	D1 – Entry of multiple pick-up points for better geographical matching
	D2 – Map-based search accentuating local area opportunities
Communicating	D3 – Standardizing seller-buyer communication for efficient interaction
Hand-over	D4 – Fast, flexible, and environmentally friendly delivery service in local area
	D5 – Custom packaging for convenience and personalization
Feedback	D6 – Feedback through emoticons on different aspects of the transaction

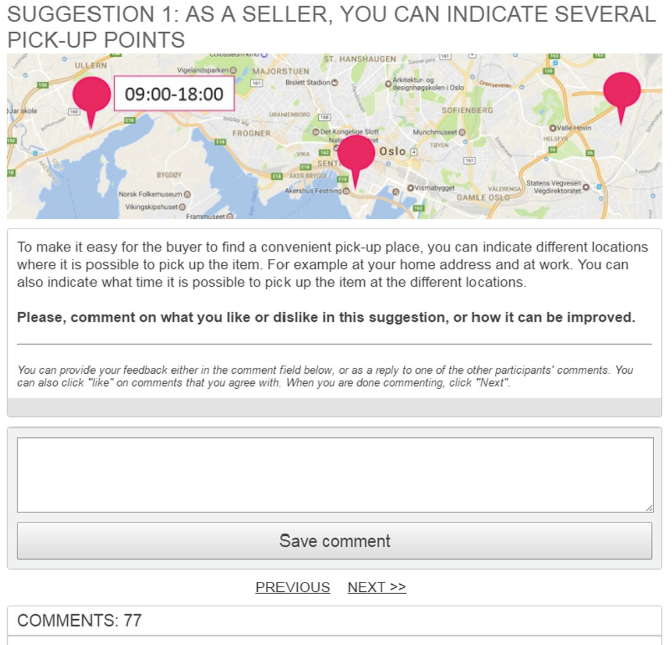

Fig. 1. Design concept (D1) in the online environment.

comments in discussion threads. One discussion thread was provided for each concept. The participants could choose to make stand-alone comments, or make their comments as replies to other participants' comments. The participant feedback was monitored by two moderators. The role of the moderators was to follow up on the comments and engage users in discussions, through positive acknowledgement of constructive comments and follow-up questions to gather more in-depth understanding of a comment.

Whenever a participant received a reply on one of their comments, the participant received an email notification with a link to the discussion thread to reply. Previous studies have shown that the presence of a moderator in asynchronous online feedback session increases the quality of the participant feedback [14].

In total, 103 participants were recruited to the study through a market research company. Of these, 89 provided one or more comments as feedback. The sample was balanced on gender, and mean age was 39 years (SD = 13, min = 18, max = 74). All participants had experience in online redistribution markets. The study was conducted in English, and the feedback period lasted for two weeks.

4.4 Users Design Feedback

Users' feedback to the design concepts was overall rich, providing comments to all designs. After the end of the two-week period, the study moderators extracted the data from the online environment, removing person identifiable information (e.g. email addresses). The two moderators independently coded the comments in a content analysis, applying three main codes: positive comments, comments containing negative issues or problems, and suggestions. This coding approach has previously been shown useful as a starting point for analysis of users' design feedback [14]. The interrater agreement indicated good reliability in the coding (Cohens kappa 0.6–0.8 for the different categories). Subsequently, for each of the three main coding groups, the analysts identified themes exploring the strong and weak aspects of the concepts, as well as the concrete suggestions for improvement.

An overview of the participant feedback is presented in Table 2, illustrating how the design feedback was useful to the team. First, the design feedback supported prioritizing of the concepts, second, the design feedback served to identify strong and weak aspects of the concepts, and third, the design feedback provided suggestions as to how to move forward in the design process.

Table 2. Overview of participant feedback to the design concepts

	D1	D2	D3	D4	D5	D6	Sum
Positive	40	33	50	32	5	18	180
Negative/problem	16	21	6	34	44	24	144
Suggestion	13	11	17	3	5	8	57

Prioritizing concepts. The users design feedback provided a useful basis for differentiating between the concepts with regard to which ones to take forward. It may be noted that the mainly positive feedback for design concepts 1 and 3, as observed in Table 2, suggests that these resonate with a perceived user need, whereas the rest of the concepts raised mixed comments and many concerns about their implementation. In particular, design 3 should be prioritized in further development. The participant comments served to underscore and substantiate reasons for likes and dislikes, as exemplified in the comments below.

> *Very good idea! This will make it easier for both parts, and shorten the discussion back and forth.* (Participant comment, D3)

> *I do think the emoticons makes it look kinda childish.* (Participant comment, D6)

Identifying strong and weak aspects in the concepts. The design feedback also provided useful insight into the detailed strong and weak aspects of the concepts. This is useful in the subsequent design process, when deciding how to refine or move forward with the prioritized concepts. Such aspects could concern context of use, assumptions concerning the target users and possible future use, as well as design issues, as exemplified below.

> *I like this suggestion as it is. To me, a lot of people seem to assume that everyone has a car available for pick up, which is not the case at all. Also, a lot of people get too locked on the idea that they need to do their deal from home.* (Participant comment, D1)

> *I do not like to give too much information in the ad about who I am and where I live. I do not want my name/address to come up in a google search. I prefer to give more information when I have a buyer.* (Participant comment, D1)

Gathering suggestions to drive the subsequent design process. Finally, the participant comments served to gather suggestions on how to move forward in the design process. Such comments are particularly useful for the designers, as they may serve as means to get needed creative input in the design process. Interestingly, it may be noted that the design concepts receiving the most suggestions were also the ones receiving the most positive comments, indicating that suggestions are driven by participants' enthusiasm for the concepts.

The suggestions concerned specific changes to the concept, as well as suggestions on which contexts the concepts would be particularly suitable for, as exemplified in the following comments.

> *Should be possible to save your work address, in addition to your home address, in your profile [...]. Then both could automatically be suggested as regular pick up places.* (Participant comment, D1)

> *Good idea, but not too many pick-up points at same time. Better with a neutral pick-up point, I don't like to give my home address to everyone.* (Participant comment, D1)

5 Lessons Learnt and Way Forward

5.1 Lessons Learnt

The presented case is interesting in that it demonstrates the viability for user feedback on early design concepts even for services as complex as those of the sharing economy. While practical user trials of such concepts are challenging in an early phase, if not impossible, due to the complexity and networked character of sharing economy services, it is highly useful to see that users are able to provide qualified feedback on concepts, even on the basis of just an illustration and a high-level description.

While the applied approach to online design feedback was found useful for differentiating between concepts, the qualitative character of the feedback also enabled the exploration of strong and weak aspects of the concepts as well as concrete input on how to move forward with the design. As such, this means of user involvement may be a valuable complement to the design process.

The approach of presenting the design concepts in the form of digital idea cards with an illustration and a high-level description text was useful in the sense that it allowed feedback from a broad range of users, while at the same time enabling differentiation between concepts and relevant input to design. It may in particular be noted that the more positively perceived concepts were also the ones gathering more design suggestions. This indicates the need to represent the concept in text and illustrations that are not only clear and precise, but also that serve to engage the participants by highlighting issues assumed particularly valuable for the user.

The format of the idea cards, with their relatively low level of detail, may also encourage participants to keep an open mind to challenges and opportunities in the concept. Possibly more so than if highly detailed concepts or prototypes are presented.

The digital idea card approach also enables participants to provide feedback on a range of concepts, as the low level of detail does not require too much time for assessing a single concept. Hence, this approach allows the exploration of concepts covering the entire transaction process, rather than just a single stage.

5.2 Future Challenges and the Way Forward

Sharing economy services are complex and in constant change due to a number of factors. The rapid advances in information technology in our digital age leads to a perpetual emergence of new communication channels. New business models emerge as service providers outsource parts of their service systems, forming service delivery networks [16]. Digital platform providers may engage in multi-sided markets for example by adding financial or logistics partners. From a business perspective, turning collaborative consumption services into multi-sided markets may increase efficiency and reduce costs. However, in such a changing landscape it is challenging to provide coherent and flawless service experiences. Interest in service research is spreading globally, and methods to measure and improve service experience has been identified as a research priority [17]. Finding efficient ways to engage users in early stages of service design is thus of high importance.

While the presented approach seems promising for efficiently involving users in early phase design of sharing economy services, the complexity of the sharing economy likely requires a broader range of practices for user involvement in innovation that what is presented here. One possible direction, given the recent advances in artificial intelligence and automated agents such as chatbots, could be to explore the use of such technologies for user research and involvement. For example, to consider how online feedback session could be improved by the use of moderating chatbots, or how user input could be analysed through machine learning approaches, where user insight could be abstracted from much larger datasets than what are practically feasible to analyse manually. Taking advantage of technological advances in user involvement activities may support innovation in rapidly evolving areas such as the sharing economy.

Acknowledgement. The presented work is part of the Conserve & Consume project, funded by the Research Council of Norway (project no. 235526).

References

1. PwC Press Room. http://press.pwc.com/News-releases/europe-s-five-key-sharing-economy-sectors-could-deliver–570-billion-by-2025/s/45858e92-e1a7-4466-a011-a7f6b9bb488f. Accessed 01 May 2017
2. European commission report: a European agenda for the collaborative economy, COM 356 final (2016)
3. Issa, T., Isaias, P.: Sustainable Design: HCI, Usability and Environmental Concerns. Springer, London (2015). doi:10.1007/978-1-4471-6753-2
4. Belk, R.: You are what you can access: sharing and collaborative consumption online. J. Bus. Res. **67**(8), 1595–1600 (2014)
5. Botsman, R., Rogers, R.: What's Mine is Yours: How Collaborative Consumption is Changing the Way We Live. HarperCollins, London (2011)
6. Rochet, J.C., Tirole, J.: Two-sided markets: a progress report. RAND J. Econ. **37**(2), 645–667 (2006)
7. Kujala, S.: User involvement: a review of the benefits and challenges. Behav. Inf. Technol. **22**(1), 1–16 (2003)
8. Blomkvist, J.: Representing future situations of service: prototyping in service design. PhD thesis, Dissertation No. 618, Linköping University (2014)
9. Polaine, A., Løvlie, L., Reason, B.: Service Design: From Implementation to Practice. Rosenfeld Media, New York (2013)
10. Knutsen, J., Morrison, A.: Have you heard this? Designing mobile social software. FORMakademisk **3**(1), 57–79 (2010)
11. Maguire, M.: Methods to support human-centred design. Int. J. Hum Comput Stud. **55**(4), 587–634 (2001)
12. Reyes, L.F.M., Finken, S.: Social media as a platform for participatory design. In Proceedings of the 12th Participatory Design Conference, pp. 89–92. ACM Digital Library (2012)
13. Näkki, P.: Service co-design using online ideation and face-to-face testing: case city adventure. In Proceedings of the Service Design and Service Innovation Conference - ServDes. 2012, pp. 177–188. Linköping University Electronic Press (2013)
14. Følstad, A., Haugstveit, I.M., Kvale, K., Karahasanovic, A.: Design feedback from users through an online social platform: benefits and limitations. Interact. Comput. **28**(4), 421–436 (2016)
15. Yin, R.K.: Case study research: design and methods. SAGE publications (2013)
16. Tax, S.S., McCutcheon, D., Wilkinson, I.F.: The service delivery network (SDN): a customer-centric perspective of the customer journey. J. Serv. Res. **16**(4), 454–470 (2013)
17. Ostrom, A.L., Parasuraman, A., Bowen, D.E., Patrício, L., Voss, C.A., Lemon, K.: Service research priorities in a rapidly changing context. J. Serv. Res. **18**(2), 127–159 (2015)

A Qualitative Methodology for the Validation of a Common Information Space to Improve Crisis Management: Results from the SecInCoRe Project

Simona De Rosa(✉)

T6 Ecosystems, Rome, Italy
s.derosa@t-6.it

Abstract. Since the 1992 Earth Summit, 4.4 billion people or 64% of the world's population have been affected by disasters and the number of crisis has more than doubled [8, 15]. In a world where disasters strike more frequently and with more intensity, governments and emergency services have a strong incentive to increase efficiency and resilience of information systems used to plan or to manage the crises. The SecInCoRe project, financed by the European Union under the FP7 Security framework, draws a socio-technical concept of a Common Information Space (CIS) to solve the most urgent challenges faced by emergency services and first responders. Among others: enhancing interoperability and communication within different organisations at national and at European level, promoting the involvement and participation of the users in line with what is requested by Next Generation Internet Initiative. In parallel to the CIS concept design, the project developed a strategy for the validation and evaluation of the project components and concepts through the use of a CIS Demonstrator. The SecInCoRe Validation and Evaluation Strategy (VES) is based on a combined approach which integrates elements of the E-OCVM methodology [6] with the SEQUOIA methodology [11]. This paper presents the strategy that was used to validate and evaluate the SecInCoRe outcomes and its results. The aim is to report on the feedback of users from emergency services to make clear how and why a socio-technical system as the CIS could help them in the management of a crisis.

Keywords: Security · Common Information Space · Qualitative methodology · European project

1 Introduction

Nowadays, Next Generation Internet (NGI) is a cornerstone of reflections and debates about the relation between technologies and society. According to the European Commission, NGI and Internet of the Future "should provide better services, more intelligence, greater involvement and participation[1]". The aim of NGI is the improvement of socio-economic conditions of all societies worldwide. To improve socio-economic conditions and express its potentialities, NGI has to be designed for humans. In line with this, the final aim of the European Commission is "to shape this Future Internet as a powerful, open, data-driven, user-centric, interoperable platform ecosystem[2]".

[1] Available at https://ec.europa.eu/digital-single-market/en/next-generation-internet-initiative.
[2] Available at https://ec.europa.eu/digital-single-market/en/next-generation-internet-initiative.

I. Kompatsiaris et al. (Eds.): INSCI 2017, LNCS 10673, pp. 231–241, 2017.
https://doi.org/10.1007/978-3-319-70284-1_19

Cloud computing has to be considered as one of the most influential trends that are effecting the Internet of Today and the Internet of the Future, creating an interoperable world where devices, people, data and process are tremendously connected. According to Taylor [14] “the cloud is fundamentally changing the way that technology is delivered, giving rise to a host of new, previously unimaginable services. Simply put, cloud is the practice of using a network of remote servers hosted on the Internet - rather than a local server - to store, manage, and process data connected over fast and increasingly mobile Internet connections”. The cloud, indeed, allows users to access contents anywhere, anytime, on any device, over any network, improving the chance to share information and access data making networked connections the central point of citizens’ lives. Among other advantages, cloud based services can increase accessibility, instant scalability and with this ensuring a high level of interoperability [1]. Due to such positive effects, cloud computing can also be used in crisis management in order to reduce the most critical issues faced in the Security field: communication and interoperability barriers [13].

Disasters, indeed, are becoming always more frequent creating huge damages and losses [9] that in some cases could also depend from inadequate tools and procedures adopted before and during the crises. This implies the high need to create new Collaborative Crisis Management tools. In line with this, cloud based service has been selected as the starting point to build up a Common Information Space (CIS) concept in order to support first responders in emergency management. The concept has been developed by the SecInCoRe (Secure Dynamic Cloud for Information, Communication and Resource Interoperability based on Pan-European Disaster Inventory) project.

SecInCoRe was a EU-funded project that aimed at fostering interoperability and information exchange at national and European level during all steps of the crisis management cycle to reply to current challenges faced in the Security field.

Starting from the description of SecInCoRe, main project outcomes are reported in Sect. 2. In Sect. 3 the SecInCoRe methodology to validate the project’ components and to evaluate the project overall impact is described. Section 4 reports the main results from validation and evaluation activities. The paper ends with conclusions on how a qualitative approach can be applied to technological development and why this is relevant when reflecting on NGI and Internet of the Future.

2 SecInCoRe at a Glance

The concept of CIS has been explored since the 90s within the Computer Support Collaborative Work (CSCW) field, providing valuable insights on the relation between information systems and practitioners [2] and on the collaborative nature of the technology.

Within this debate, SecInCoRe has recently identified the CIS as a promising solution to enhance cross border collaboration and information exchange in crisis management, designing a dynamic and secure cloud based CIS, as a socio-technical concept that facilitates information exchange for responders and practitioners that intend to share information in a secure, useful and trusted manner. Indeed, the CIS concept based on a cloud computing approach has been considered by the project as the starting point of next generation, secure cloud-based pan-European information management in disaster preparedness, response and recovery.

As stated by Kuhnert [7] the idea behind the concept is that adopting the right technical support, people can build up a collaborative space in which information can be more securely shared and managed. In this sense, the CIS has been designed as a socio-technical system, co-designed and co-developed between end users, researchers and designers fostering and promoting the user centric approach.

A collaborative design methodology has been crucial in the concept design as well in the validation phase, helping to collect "into one conversation multiple perspectives, forms of expertise, and contexts, as it explores the interplay between the social, technological, and organizational through hands-on engagement with prototypes" [12]. For the entire process the aim has been to translate users' requirements into a working CIS demonstrator, indeed, as stated by Petersen [12] "participants become a collective resource for design and produce an environment of mutual learning".

Through the use of a co-design methodology, the SecInCoRe CIS concept[3] has been developed, as reported in Fig. 1, according to three pillars: CIS Specification, Reference Implementations and Demonstrations and Pilot Studies. Within each pillar the project has identified the components that should be integrated, both at a conceptual and technical level, to create a functioning CIS.

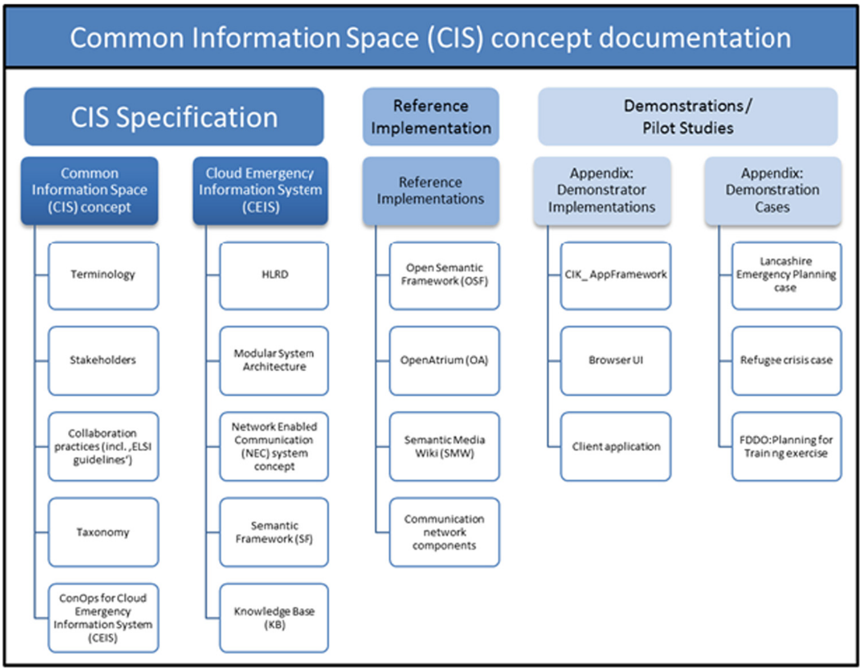

Fig. 1. CIS concept design [4] (Available at http://www.secincore.eu/wp-content/uploads/2017/05/D5.5_Evaluation-and-validation-report_V4_FIN1.pdf).

[3] Deeper information is available at www.secincore.eu.

From a conceptual point of view, the CIS specification reports l documentation to be used to create a CIS. The Reference Implementations provide the translation of the CIS Specification into a living system. These have been used during the project to build the CIS demonstration for validation purposes. Then, the Demonstration/Pilot Studies contains all information about the demonstrator implementation used to build up demonstration cases.

On the other hand, the technical implementation is based on components aimed at providing a modular concept that is flexible and adaptable to users' needs. The components developed are:

- a Network Enabled Communication (NEC) concept enabling seamless, robust and role-based access to the CIS;
- a Pan-European inventory of data sets, processes, information systems, business models used by first responders and police authorities;
- a Knowledge Base to translate parts of the inventory in a technical representation and also interlinked the content like several databases with a semantic model, the implemented SecInCoRe taxonomy;
- the Semantic search bridging between different mental models and enable easy information retrieve by using an ontology-driven search functionality [12].
- the Ethical, Legal and Social Issues (ELSI) guidance[4] has been developed to support the understanding of differences and similarities of the various involved organisations and to help users to improve the use of technology for disaster risk management.

3 A Qualitative Approach for the Validation and Evaluation of SecInCoRe

According to the components described in the previous paragraph, a strategy for the validation and the evaluation of the outcomes has been formulated as a crucial part of the project to understand if the CIS could improve current procedures for crisis management and to evaluate its potential impact.

In order to guarantee a comprehensive validation of the entire process, the SecInCoRe Validation and Evaluation Strategy (VES) has been structured on the combination of two different methodologies. Elements of the European Operational Concept Validation Methodology (E-OCVM) methodology for the direct validation of technical outcomes have been integrated with elements inspired to the SEQUOIA methodology for the evaluation of expected impacts (Fig. 2).

Concerning the validation approach, the strategy is inspired by the E-OCVM case-based approach being centered on Demonstration Cases. The aim of the approach is that during a Demonstration Case, end-users are introduced to and interact with selected elements of the SecInCoRe concept and components in a structured manner that permits the systematic collection of data and their comparison and aggregation across different Demonstration Cases [3]. The validation strategy follows a multiple-case embedded design [16] based on the aggregation of evidence collected from Use

[4] Deeper information is avialable at http://www.isitethical.eu/elsi-guidance.

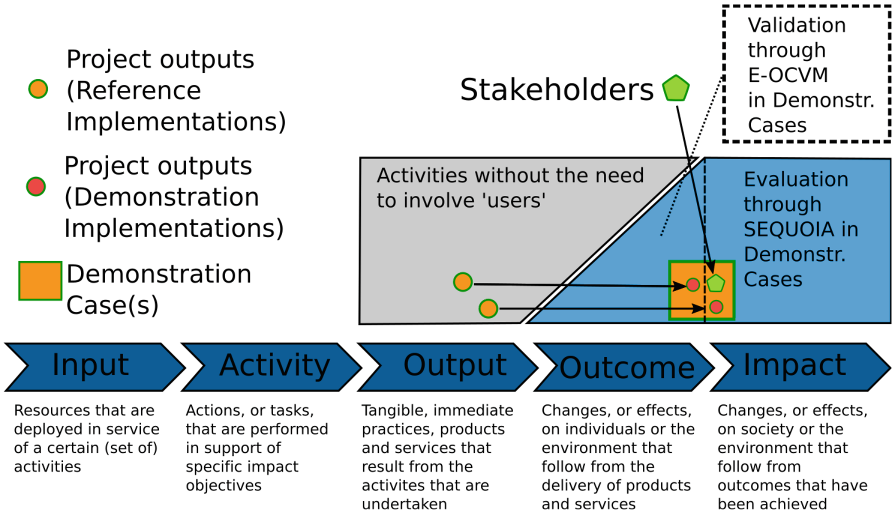

Fig. 2. The impact value chain, demonstration cases and validation/evaluation strategy [3].

Cases embedded in multiple Demonstration Cases. Demonstration Cases are based on the deployment of Demonstrator Implementations (i.e. working prototypes of different elements of the SecInCoRe concept created through the adaptation of Reference Implementations) and non-technical representations of the SecInCoRe concept (such as story telling, videos and mock-ups) in a workshop setting organised around a realistic scenario co-designed with end-users to reflect their interests and practices. As described by Cucco [5], Demonstration Cases are organised in a way that:

- permits end-users to experience the potentialities of a CIS designed according to SecInCoRe core principles;
- allows them to perform guided practical activities (Use Cases) that entail their interaction with the available technical implementations of different elements of the SecInCoRe concept.

This means that the case-based approach adopted is based on the creation of multiple Demonstrator Implementations and on the adaptation of Reference Implementations to the specific scenarios and use cases designed with the involvement of end-users for each Demonstration Case.

From a methodological perspective, the structuring and analysis of a Demonstration Case has been based on a Demonstrator Case Protocol (DCP). The DCP aims at providing guidelines and a pre-defined template, the Demonstration Case Template, to ensure that the activities are performed in a consistent way so as to allow comparability across different Demonstration Cases. The criteria used for the validation of projects' components and concepts have been derived from qualitative and quantitative SMART indicators. The indicators for each validation module are derived from the High-Level Requirements (HLR) elaborated within the project. Starting from the HLR identified to design the CIS, and the related components, specific tools (among others: structured observation, interviews, questionnaires) have been used to understand if the users agreed with the HLR and with the translation in the technical Demonstrator.

On the other hand, a structured approach to perform the evaluation of the project in order to estimate the socio-economic impact was needed as well. According to the definition provided by the International Association for Impact Assessment (IAIA), impact is "*the difference between what would happen with the action and what would happen without it*[5]". The approach adopted for the evaluation is based on the SEQUOIA methodology to assess the potential socio-economic impact [11]. The SEQUOIA methodology is based on the following five steps:

- Baseline identification;
- Mapping the areas of impact;
- Ex-post scenario description;
- Ex-post scenario quantification:
  - economic impact
  - social impact
- Final assessment analysis: calculation of iROI, xROI, tROI and RORI.

As described by Passani [10], the first step of the SEQUOIA methodology is mainly aimed at identifying project outputs that could have some improvements to the social/economic welfare and that could affect stakeholders. To do that, the basis of the methodology is the comparison between the situation without the outputs of the project and the situation modified according to the outputs of the project. In line with this, the description of a situation without the project output is defined as the baseline scenario or, in other words, as the zero scenario. The second step is focused on the identification of areas of impact in order to analyse which kind of socio-economic value is produced by the project. This third step aims at describing/forecasting the situation after the completion of the project and after the exploitation of the research outputs. The fourth step is related to the quantification of the ex-post scenario and led to transform qualitative information in quantitative data that could describe the economic and the social impact of the project. Finally, the final step of this methodology relates to the analysis of all collected data.

As said, the basis of the methodology is the assessment of the impact of a project through the adoption of counterfactual methods based on the comparison of so-called *zero scenarios* (the current practice) with a future scenario in which SecInCoRe has been adopted. Due to the fact that during the project it was not possible to deploy a fully functional system based on the SecInCoRe principles and concept in a real-life setting, the possibility to perform a counterfactual analysis was limited. In line with this, the assessment has followed a framework based on the following steps:

- Development of narratives (Evaluation Scenario) co-generated with end-users, aimed at identifying a realistic operational situation that provides the background for the challenges that end-users involved in a Demonstration Case are facing;
- Understanding with end-users how SecInCoRe could help addressing the challenges posed by the Evaluation Scenario;

[5] Available at http://www.iaia.org/uploads/pdf/What_is_IA_web.pdf.

- Let the stakeholders compare the Evaluation Scenario with SecInCoRe to the Evaluation Scenario without SecInCoRe and assess the perceived benefits of SecInCoRe across a series of relevant impact areas and dimensions identified on the basis of the SEQUOIA methodology.

According to the SEQUOIA methodology four main dimensions related to the social impact have been analysed in SecInCoRe:

- Impact of the CIS on working routine and on current procedures;
- Impact on knowledge and information sharing;
- Impact on networking;
- Impact on social capital increment for users and participants.

To apply the above described VES methodology, the process has been developed in order to reach the following aims [5]:

- coordinate the activities of different teams and facilitate the flow of information between them;
- guide the design of preparatory activities;
- check the alignment of Demonstration Cases with SecInCoRe's VES objectives;
- ensure the collection of background data according to a standardized format.

Regarding the data collection it is fair to say that different techniques have been used selecting the most appropriate tool for the kind of activity performed. As for the validation, structured observations, questionnaires, semi-structured interviews and focus groups have been selected in order to maximise the opportunity to collect data. In addition, a standard and structured questionnaire has been developed and distributed to all participants involved in the validation and evaluation activities to collect standardized information on the background of each participant, their current job and their previous experiences with CIS in crisis management [5].

Once ended the data collection, the analysis of data and evidences collected started through two coding cycles. The first coding cycle was based on an a-priori list of codes derived from the High-Level Requirements related to the elements of the SecInCoRe concept used in the Demonstration or Pilot Demonstration Case. The first coding was individually undertaken by each observer/interviewer; results were then compared and triangulated. The aim of the second coding cycle was twofold: (a) generate aggregate categories from the individual codes; (b) identify emerging themes or topics that were not originally included in the a priori list.

Results coming from the analysis are described in the next paragraph.

4 Main Results from the Validation and Evaluation

Based on the methodologies introduced above, three validation activities and three evaluation activities have been performed. The validation and evaluation activities have seen the involvement of 39 users. Out of the 39 users, 25 are first responders engaged in different fields (e.g., Fire Brigades, Civil Protection, Police, Local Authorities). Most of them have been involved as part of the Advisory Board of the project and so very well aware of project development. The remaining 14 are experts in Ethical, Legal and Social Issues, which helped in the evaluation of those crucial aspects for the creation of the CIS.

Through the engagement of external users involved in the Security field it has been possible to reply to the research questions that drove the VES during the project life-time.

The validation strategy was mainly directed to reply to the question: are we designing the right system?

Validation results showed the positive feedback received by the users on the need to have a CIS, as designed by SecInCoRe, to harmonise procedures for preparedness and planning across Europe. According to the users' comments, a CIS could help to standardise practices for sharing documents and enhancing collaboration, which is something useful but above all needed. As reported by Cucco [4], "The CIS has also been confirmed as a powerful tool to keep in touch with practitioners in other countries, thanks to the chance to contact directly the author of the information stored in the CIS. This facilitates the establishment of new partnership at the European level, enlarging the source of information as well as the network of people that is possible to reach". In detail, users validated collaborative functions, such as chats, discussions boards and functions for sharing contents in real time, the graph view and the design of the Knowledge Base and the Inventory. The collected feedback underlines that collaborative functions added in the CIS could improve working routines considerably, facilitating the collaboration among different emergency services. The graph view developed by the project under the work on the semantic search has been perceived by the users as a new way to approach the search of topics when looking for documents to prepare a plan or an exercise, increasing the chance to derive to new knowledge. Finally, the design of the Knowledge Base and the Inventory has been judged as positive development to overcome the problems related to the fact that practitioners in Europe are using several different inventories which do not allow users to have a homogeneous reference to common documents.

From an evaluation perspective, it has been possible to answer two main research questions related to the impact assessment:

1. Why is SecInCoRe relevant and for whom?
2. What is the difference the SecInCoRe project makes?

From the evaluation activities emerged that "SecInCoRe is relevant because it has produced the design of a CIS that, according to the stakeholders, can improve current practices on preparedness and planning in the emergency services, effectively addressing some of the major issues that still emerge from current practises such as: the lack of information exchange, the difficulties in creating new partnership and the obstacles to collaborate in an easy and effective way" [4]. Due to its scope, the project is relevant mainly for practitioners and first responders who regularly work in the planning and preparedness phase of the emergency at the European as well as at the national level.

Regarding the kind of difference that SecInCoRe can make, it is fair to say that even if the original methodology identified several potential areas of impact, the social impact was the most relevant area to assess for the CIS designed by SecInCoRe. In line with this, the social impact is assessable looking at the already mentioned variables, described in detail in the next paragraphs below.

Impact of the CIS on working routine and on current procedures
Results of the investigation show that the use of a CIS designed according to SecInCoRe principles could improve the current procedures used for crisis management. Currently, practitioners around Europe do not use a unique tool for accessing and exchanging information or to collaborative produce a plan. In line with this, "establishing a CIS with access for the practitioners will allow to get access to the same critical mass of information, having access to documents and material and contextually using the collaborative functions to set up collaborative groups" [4]. A quantitative estimation of the time saved was not possible to derive, however stakeholders agree that such system could really help in saving time to access and share the information.

Impact on knowledge and information sharing
As already mentioned, it emerged clearly that there are currently no tools or repositories at the European level that can easily support the practitioners in finding information during the planning phase. Therefore most of the stakeholders appreciated the opportunity to have access to the same critical mass of information embedded in a CIS where different emergency services across Europe can contribute. Indeed, "having a common repository could improve the knowledge of sources used by other practitioners and promote the standardisation of the documentation, stimulating the standardisation of the produced plans" [4].

Impact on networking
An additional dimension for mapping the SecInCoRe's impact is linked to the opportunity to establish new partnerships among practitioners engaged in the emergency services. A current issue faced by emergency services is related to the fact that, generally, the network of people contacted to receive additional information is generally limited to the personal contacts that each practitioner has. On the contrary, the SecInCoRe CIS could offer the opportunity to build new relations with colleagues from other countries even outside from personal relationships. In detail, the possibility to contact the author of the information stored in the inventory, as foreseen by the system, would allow to directly access the sources of relevant information and with this extremely enhancing collaboration.

Impact on social capital increment for users and participants
Regarding the impact on the dimension related to social capital among users of the same CIS, the most important element that emerged is about trust. Due to the high sensibility of the work, practitioners are very concerned about trust issues. In line with this "a CIS with a high-level managing authority that could guarantee the process and the management of the system would be extremely important to guarantee the establishment of a trusted network. On such bases, stakeholders would be inclined to trust the system, the information stored in it as well as the security infrastructure behind it. These points would encourage subscriptions to the system and its use on a daily basis" [4]. It is fair to say that the current dimension related to the social capital differs to the previous one related to the networking due to fact that impact on the networking was mainly observed in relation to the chance to establish new contacts with other first

responders and to create potential new collaborations without exploring benefits related to the trusted circle created by the CIS, which is the case of this last dimension, observed specifically to understand the relation between CIS and trust.

5 Conclusions

This paper introduces the qualitative methodology used during the SecInCoRe project to validate and evaluate project's outcomes and summarises main results from the validation and evaluation activities. The EU-funded project delivered a CIS concept and additional components to foster interoperability and collaboration within the emergency services. The methodology used to understand if a CIS could be a valuable concept to be used in the Security field has been based on the integration of a Demonstration Case approach with a methodology for the impact assessment in order to gather a broader perspective from the users engaged in the validation and evaluation process. Main results show that a SecInCoRe CIS could harmonise the tools and methods used at European level by supporting the use of the same functionalities and fostering collaboration at a wider level among emergency services. In addition, stakeholders recognised the capacity of the system to address some major limitations of the systems currently used for crisis management for instance improving collaboration practices in Europe.

It is fair to say that due to the fact that investigations have been conducted mainly through the use of a qualitative methodology, it was not possible to quantify the difference that SecInCoRe could produce in terms of economic impact. However, through the potentiality of the qualitative approach is possible to trace the capacity of the consortium to create a CIS concept based on a user centric approach. Indeed, feedback from the users has been constantly considered in further versions of the CIS Demonstrator providing always an updated version of the project outcomes based on users' experience. Engaging frequently users and stakeholders made it possible to deliver solutions created for the humans and so they are connected to their needs and requirements. The participation and the engagement of first responders and people highly involved in crisis management has been the key of the methodology. To conclude, the CIS concept as designed by SecInCoRe has been judged by the users as a tool that could strongly modify current practices and routines used in the emergency services, solving some major issues related to the lack of interoperability and information sharing. The CIS based on cloud service could genuinely change the way in which crisis are managed, improving both the preparedness and training phase as well as the response phase, also thanks to the use of cloud based services responding to the current and urgent emergent security issues with a solution created by the humans for the humans.

Acknowledgements. The research leading to these results has received funding from the European Union Seventh Framework Program (FP7/2007-2013) under grant agreement no 607832 (project SecInCoRe). The text reflects the authors' views. The European Commission is not liable for any use that may be made of the information contained therein. For further information see http://www.secincore.eu/.

References

1. Armbrust, M., Fox, A., Griffith, R., Joseph, A.D., Katz, R.H., Konwinski, A., Zaharia, M.: Above the clouds: a Berkeley view of cloud. In: Electrical Engineering and Computer Sciences. University of California at Berkeley (2009)
2. Bannon, L., Bødker, S.: Constructing common information spaces. In: Bannon, L., Bødker, S. (eds.) Cooperative Work, pp. 81–96. Springer, Dordrecht (1997). doi:10.1007/978-94-015-7372-6_6
3. Cucco, I., De Rosa, S., Grey, P., Despotov, B., Schäfer, C., Sauerland, T., Petersen, K., Danilidis, I., Hirst, P.: SecInCoRe Deliverable 5.4 – Validation report and final evaluation model of communication system concept (2016)
4. Cucco, I., De Rosa, S., Grey, P., Despotov, B., Schäfer, C., Sauerland, T., Petersen, K., Danilidis, I., Hirst, P.: SecInCoRe Deliverable 5.5 – Evaluation and Validation Report for SecInCoRe Stakeholders (2017)
5. Cucco, I., De Rosa, S., Passani, A., Georgiev, A., Pottebaum, J., Schäfer, C., Sauerland, T., Petersen, K., Danilidis, I.: SecInCoRe Deliverable 5.3 – Validation strategy and first functional evaluation model of communication system concept (2016)
6. Eurocontrol: European Operational Concept Validation Methodology. Version 3.9, Vol. I (2010)
7. Kuhnert, M., Wietfeld, C., Paterour, O., Georgiev, A., Petersen, K., Büscher, M., Pottebaum, J.: Next generation, secure cloud-based Pan-European information system for enhanced disaster awareness. In: Palen, L., Buscher, M., Comes, T., Hughes, A. (eds) In: Proceedings of the 12th International Conference on Information Systems for Crisis Response and Management (Iscram), Kristiansand (2015)
8. Munich, R.E.: NatCatSERVICE Loss Events Worldwide 1980–2014. Munich Reinsurance, Munich (2015)
9. NatCatSERVICE: Münchener Rückversicherungs-Gesellschaft, Geo Risks Research (2012)
10. Passani, A., De Rosa, S., Georgiev, A., Buescher, M., Petersen, K.: SecInCoRe D5.2- Early setup of evaluation model for internal use cases (2015)
11. Passani, A., Monacciani, F., Van Der Graaf, S., Spagnoli, F., Bellini, F., Debicki, M., Dini, P.: SEQUOIA: a methodology for the socio-economic impact assessment of software-as-a-service and internet of services research projects. Res. Eval. **23**(2), 133–149 (2014)
12. Petersen, K., Buscher, M., Kuhnert, M., Schneider, S., Pottebaum, J.: Designing with users: co-design for innovation in emergency technologies. In: Palen, L., Buscher, M., Comes, T., Hughes, A. (eds) Proceedings of the 12th International Conference on Information Systems for Crisis Response and Management (ISCRAM), Kristiansand (2015)
13. Schäfer, C., Sauerland, T., Pottebaum, J., Marterer, R., Behnke, D., Wietfeld, C., Despotov, B.: Cloud-based semantic services for pan-European emergency preparation and planning. In: 2017 Annual IEEE International Systems Conference (SysCon), pp. 1–7. IEEE (2017)
14. Taylor, S.: The next generation of the internet revolutionizing the way we work, live, play, and learn. In: CISCO Point of View (2013)
15. UNISDR: Impacts of disasters since the 1992 Rio de Janeiro Earth Summit (2012). http://www.preventionweb.net/files/27162_infographic.pdf
16. Yin, R.K.: Case study research, design and methods, 3rd edn. Sage Publications, Newbury Park (2003). ISBN 0-7619-2553-8

Data-Driven Research and Design

Ewya: An Interoperable Fog Computing Infrastructure with RDF Stream Processing

Eugene Siow(✉), Thanassis Tiropanis, and Wendy Hall

Electronics and Computer Science, University of Southampton, Southampton, UK
{eugene.siow,t.tiropanis,wh}@soton.ac.uk

Abstract. Fog computing is an emerging technology for the Internet of Things (IoT) that aims to support processing on resource-constrained distributed nodes in between the sensors and actuators on the ground and compute clusters in the cloud. Fog Computing benefits from low latency, location awareness, mobility, wide-spread deployment and geographical distribution at the edge of the network. However, there is a need to investigate, optimise for and measure the performance, scalability and interoperability of resource-constrained Fog nodes running real-time applications and queries on streaming IoT data before we can realise these benefits. With Eywa, a novel Fog Computing infrastructure, we (1) formally define and implement a means of distribution and control of query workload with an inverse publish-subscribe and push mechanism, (2) show how data can be integrated and made interoperable through organising data as Linked Data in the Resource Description Format (RDF), (3) test if we can improve RDF Stream Processing query performance and scalability over state-of-the-art engines with our approach to query translation and distribution for a published IoT benchmark on resource-constrained nodes and (4) position Fog Computing within the Internet of the Future.

Keywords: Fog computing · Stream processing · Interoperability · Internet of Things · Query translation · Linked Data · Workload management

1 Introduction

In the science-fiction motion picture Avatar[1], Eywa is the name of a biological internet on the planet Pandora, made up of trees which are distributed over the surface of the planet, that store and process information and memories. The flora and fauna of the planet form an ubiquitous sensor network feeding Eywa.

The Internet of Things (IoT) is growing to become a similar ubiquitous network of sensors and actuators for our planet. Fog computing is an emerging technology which seeks to bridge a gap for the IoT, like the fictional Eywa on Pandora, in between the ground where sensors and actuators are deployed and

[1] http://www.avatarmovie.com/.

I. Kompatsiaris et al. (Eds.): INSCI 2017, LNCS 10673, pp. 245–265, 2017.
https://doi.org/10.1007/978-3-319-70284-1_20

collect/act on data, and the cloud [12], where larger amounts of resources for processing and storing data can be provisioned dynamically.

Some challenges for existing and past proposals of distributed system architectures within the IoT lie in (1) how the heterogeneity of device, platform and data (which is often overlooked at the architecture level) is managed, (2) how multiple streams of data can be processed in a performant, scalable way in real-time and (3) how and what resources can and should be provisioned for these real-time applications. Furthermore, the architecture should also provide a means to manage the *data plane* of the network, e.g. the streams, and the *control plane* of the network, e.g. distribution of application processing on streams.

Bonomi *et al.* [7], in their paper on Fog Computing for the IoT, introduce the defining characteristics and requirements of Fog Computing systems as (1)*interoperability* to support heterogeneous devices and data, (2) *low latency/high performance* for streaming data and real-time applications and (3) distributed infrastructures that have the ability to *scale* horizontally due to the potential wide-spread distribution and large number of resource-constrained Fog nodes.

As Fog Computing extends to the edge of the network, it benefits from low and predictable latency due to data locality, location awareness, mobility, wide-spread deployment and geographical distribution [6]. Furthermore, Fog Computing infrastructures deployed on a network of resource-constrained devices can enhance and support big data processing and analysing data in the cloud by taking advantage of data locality (therefore low latency), mobility and distribution, to provide performant, interoperable and scalable services for data integration and supporting streaming, real-time applications before cloud processing. An example Fog Computing application given by Bonomi *et al.* [6] is a Smart Traffic Light system that has local subsystem latency requirements in the order of <10 ms while requiring deep analytics over long-periods in the cloud.

The contribution of this paper is to propose a new Fog Computing architecture, Eywa, and provide the components for stream processing with it, focused on streams and data interoperability and to evaluate it based on the key metrics of interoperability, performance and scalability (measures of a Fog Computing system as defined by Bonomi *et al.* [7]) by:

1. Formally defining and implementing a means of distributing query workload in the *control plane* with an inverse publish-subscribe and push mechanism.
2. Showing how streaming data from heterogeneous devices can be integrated and made interoperable through organising data as Linked Data in the Resource Description Format (RDF) and taking advantage of the shape of time-series IoT data to optimise performance in the *data plane*.
3. Evaluating our approach and framework, by the metrics of latency and scalability, against state-of-the-art RDF Stream Processing engines using a published Smart City benchmark, CityBench [1]. We record improved performance and scalability for real-time applications with streaming queries on inexpensive, mobile, resource-constrained Fog nodes.

The rest of the paper is organised as follows: we formally define a framework for Fog Computing that allows the distribution of stream query workload for real-time applications in Sect. 2 and go on in Sect. 3 to introduce Linked Data and the RDF format and how it can be used to integrate and make interoperable input and output from heterogeneous IoT devices. Next, Sect. 4 explains RDF stream processing (RSP), how it can be optimised for time-series IoT streams and how its implemented in Eywa. We then evaluate Eywa in Sect. 5 with results and discussion in Sect. 6 based on the metrics of performance and scalability. Finally, we argue for the position of Fog Computing within the Internet of the Future as an IoT solution in Sect. 7 and discuss related work in Sect. 8.

2 Eywa: An Infrastructure for Fog Computing

The purpose and contribution of this section is to introduce and formally define Eywa, (1) an infrastructural Fog Computing framework where stream processing can be performed on resource-constrained lightweight computer nodes like Raspberry Pis' (RPis'), (2) where the processing workload can be distributed amongst nodes within a inverse-pub-sub *control plane* and (3) each node can maintain its own independence and control over access, resources, security and privacy.

An Eywa network, that forms the basis for the Eywa Fog Computing infrastructure, is explained in Definition 1. The utility of this network for processing streams is to facilitate (1) stream query delivery to relevant nodes, (2) *stream query delivery* to relevant nodes, (3) *distributed processing* and 4) *results delivery* to requesting nodes. To this end, Definition 2 defines the 3 types of nodes in the network, Sect. 2.1 explains query delivery, Sect. 2.2 distributed processing and Sect. 2.3 results delivery.

Definition 1 (Ewya Network, ε). *An Eywa network, ε, consists of a set of nodes, N and connections, C, such that $\varepsilon = (N, C)$. Each node, $n \in N$, can be a source(s), client(τ) or broker(b) node, such that $n = \{s, \tau, b\}$. Each connection, $c \in C$ exists uniquely between two nodes such that $C \subseteq N \times N$.*

Definition 2 (Source, s, Client, τ and Broker, b nodes). *Given the set of source nodes, S, client nodes, T, and broker nodes, B, where $N = S \cup T \cup B$. A source node, $s \in S$, is a node that produces a set of time-series streams, Γ. A client node, $\tau \in T$ is a node with a set of queries Q expecting a set of corresponding results, R. A broker node, $b \in B$, establishes a connection, c, for new source and client nodes to enter the network.*

As *source* and *client* nodes form up in a Eywa network, *broker* nodes (1) provide a point of entry for new nodes into the network, (2) provide a point of entry for new nodes into the network, (3) not store or process but forward data, (4) consume minimal resources and (5) employ redundancy (multiple separate instances) within the network so as not be become single points of failure. Hence, facilitating both the formation and data flow in Eywa.

2.1 Stream Query Delivery by Inverse-Publish-Subscribe

Once an Eywa network has been formed, stream processing can take place as clients (τ) issue stream queries (Q). Traditionally, source nodes (s) publish data while client nodes subscribe to data. However, in Eywa, it is desirable for clients to collaborate with the sources to share the workload, hence, we propose an inverse-publish-subscribe mechanism for query delivery.

Each *source* node subscribes to a topic for each of its streams. *Client* nodes then publish queries to the relevant topics. Uniform Resource Identifiers (URIs), proven to work for the web, are used to provide a means of uniquely identifying and exchanging topic names. Definition 3 formally describes the mechanism.

Definition 3 (Inverse-publish-subscribe). *Given the set of topics, M, a source node, s, subscribes to a topic, $\mu \in M$ for each stream within Γ to form $\bigcup_{\mu \in M_\Gamma} sub(\mu)$, where M_Γ is the set of all topics of s and $sub(\mu)$ is a function that produces a subscription to μ. For each query, $q \in Q$, from a client node, τ, a distribution function, α, builds a set of query-topic pairs $\alpha(q) = (Q_\mu, M_q)$ where Q_μ is the set of all sub-queries in q, each referencing a particular topic μ and M_q is the set of all topics referenced in q. Each sub-query, $q_\mu \in Q_\mu$ is published to it's particular topic $\mu_q \in M_q$ by the publish function of query x to topic $y, pub(x, y)$, so all publications from q is represented by $\bigcup pub(q_\mu, \mu_q)$.*

2.2 Distributed Processing

Source nodes receive the queries, perform part of the processing and deliver the results to *clients* that process the results. This forms the axis of client-source collaboration and distributing processing workload as defined in Definition 4. *Source* nodes control their own resources and response to queries. Quality of service is not in the scope of this work but can be configured for best effort, service-level agreements, trustless networks or consensus protocols.

Definition 4 (Distributed Source Node Processing). *A source node, s, receives a query, q, for a topic, μ, and converts it into a work function, ω with a conversion function, $\lambda(q) \rightarrow \omega$. The work function, ω is applied to the corresponding stream for topic μ, γ_μ, where $\gamma_\mu \in \Gamma$, so that $\omega(\gamma_\mu) = \gamma_r$ and the resulting stream, γ_r, is pushed to the requesting client, τ.*

2.3 Push Results Delivery and Sequence Diagram

Client nodes receive results streams via a direct push from *source* nodes. Operations involving multiple streams, like aggregations are performed on the clients and results of the queries are published to topics as output to applications. Definition 5 details the process of results delivery and query output.

Definition 5 (Push Results Delivery). *A client node, τ, receives a set of result streams, Γ_r, by push delivery. For each query, $q \in Q$, in the set of queries for that client, a work function, ω_τ is produced by $\lambda_\tau(q) \rightarrow \omega_\tau$ and executed on $\bigcup \gamma_r$, where γ_r are all the result streams corresponding to the query, q. The result, γ_q, from $\omega_\tau(\bigcup \gamma_r) = \gamma_q$ is published to a client results topic, μ_τ.*

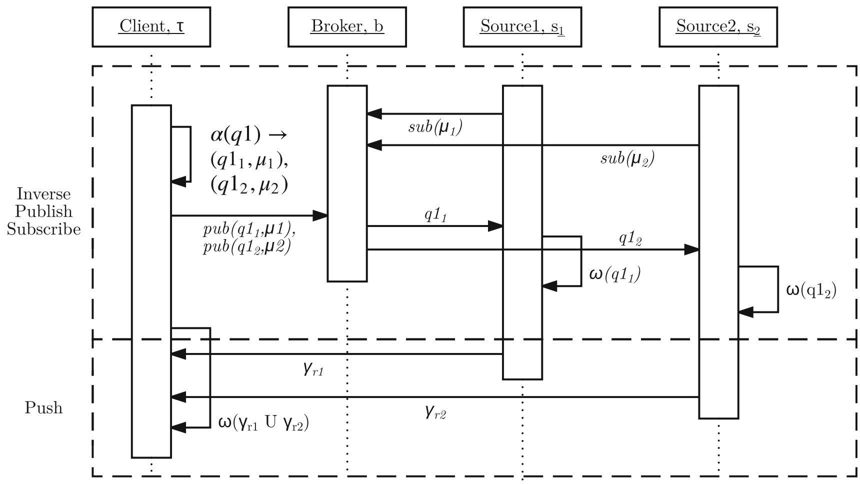

Fig. 1. Sequence diagram of stream query processing in Eywa

Table 1. Glossary of special symbols and their definitions

	Description	Definition(s)
τ	Client node	2
μ	Topic (Uniformed Resource Identifier)	3
γ, Γ	Stream, set of streams	3, 4, 5
$\alpha()$	Distribution function, build query-topic pairs	4
$\lambda()$	Conversion function, converts query to work function	4, 5
$\omega()$	Work function, applied to streams	4, 5

Figure 1 shows the full sequence diagram of the query processing process beginning from the *source* nodes subscribing to topic URIs, receiving queries when published by *client* nodes, distributed processing and result delivery. Table 1 shows a glossary of symbols used and their corresponding definitions.

3 Linked Data for Interoperability

The Internet of Things (IoT) is currently beset by product silos and to unlock its potential, an open ecosystem based upon open standards for identification, discovery and interoperation of services is required[2]. As previously noted, interoperability is an important metric for Eywa. To provide for interoperability, we consider Linked Data, a means of publishing data on the Web so that distributed structured data can be interconnected, exchanged and retrieved through semantic queries [5]. A common way of representing Linked Data is with the Resource

[2] https://www.w3.org/2014/12/wot-ig-charter.html.

Description Format (RDF) where the use and referencing of common identifiers and ontologies helps to integrate data. Furthermore, sensor ontologies like the Semantic Sensor Ontology (SSN) [9] have been developed for Linked Data which encourages providing machine-interpretable descriptions within RDF to describe what data represents, where it originates from, how it can be related to its surroundings, who is providing it, and what its attributes are e.g. a unit of measure for each sensor reading, its sensor platform and location. Barnaghi *et al.* [3] further support the view that Linked Data is a means of connecting and integrating cyber, social or physical world data in the IoT.

RDF documents are composed of tuples of triples formed from a subject, predicate and object. For example, in the statement 'sensor1 has weatherObservation1', the subject is *sensor1*, the predicate is *has* and the object is *weatherObservation1*. 'weatherObservation1 hasValue 30knots' is another triple and the set of these triples forms an RDF graph in which we know sensor1 had an observation with a value of 30knots. This is formally expressed in Definition 6.

Definition 6 (RDF triple, t and RDF graph, G). *t is a tuple $(s, p, o) \in (IB) \times I \times (IBL)$ where s, p and o are subject, predicate and object respectively. I, B and L are disjoint infinite sets of Internationalised Resource Identifiers, blank nodes and literals respectively. G is a set of triples such that $G = \{x : x \in t\}$.*

The use of Internationalised Resource Identifiers (IRIs), URIs that support unicode characters, like 'http://purl.oclc.org/NET/sao/hasValue', enables resources to be uniquely identified and referenced. Blank node's are anonymous resources without an IRI and literals are values like strings and dates. SPARQL is a language used for querying RDF.

Hence, Linked Data provides IoT interoperability as (1) RDF graphs allow structured data to be interconnected with IRIs providing common identifiers that can be referenced across sources, (2) SPARQL provides a means for semantically querying RDF graphs and (3) there exist many ontologies like the SSN Ontology [9] that provide a reference and model for organising sensor data.

4 RDF Stream Processing (RSP) for the IoT

Stankovic [18] explains that the IoT will increasingly be composed of "a very large number of real-time sensor data stream" and that a "given stream of data will be used in many different ways". RDF stream processing (RSP) is the area of work which enables Linked Data to be produced, transmitted and continuously queried as RDF streams[3]. RSP enables us to take advantage of the interoperability of RDF on streams while preserving the power of semantic queries. However, as Buil-Aranda *et al.* [8] have shown with Linked Data on the web, performance is an issue especially when dealing with a series of non-trivial queries.

[3] https://www.w3.org/community/rsp/.

We have previously surveyed about 20,000 unique IoT schema from public IoT data streams and discovered that a large majority of the sampled devices had *flat* schemata and *wide* schemata (99.5% and 76.3% respectively) [17]. Flat schemata have no nested layers (table-like rather than tree-like) and wide schemata have more than one property besides the timestamp.

We then looked at the ontologies for integrating time-series sensor data in RDF and observed that linked sensor data is produced from these ontologies as (1) *device metadata* like the location and specifications of sensors, (2) *observation metadata* like the units of measure and types of observation and (3) *observation data* like timestamps and actual readings. A large portion of the triples produced consisted observation metadata which was repetitive and made up of IRIs with additional 128-bit universally unique identifier strings. Our results showed across various scenarios that this data was often redundant and not used eventually.

It follows that Eywa's RSP for time-series streams can be optimised to:

1. Store *observation data* succinctly in flat and wide rows instead of graphs.
2. Abstract the small set of *device metadata* to store as RDF mappings.
3. Compress *observation metadata* as bindings and only materialise if needed.
4. Distribute part of a query to be applied on streams on a Eywa source node which can reduce the bandwidth required and share the workload.

When an RSP query is registered, it is translated with reference to the metadata expressed in mappings, part of the query is distributed to the relevant source streams and the rest of the query is applied on the subset of stream data received on the client side as a continuous query. Any additional metadata is materialised and added to the results. We walkthrough an example query from a smart city scenario published in CityBench [1], a published benchmarking suite for streaming applications on smart city data gathered from IoT sensors deployed within the city of Aarhus in Denmark from February 2014 to November 2015.

Table 2 shows a weather observation from a CityBench stream. This is a stream of actual flat and wide time-series observation data and consists of humidity, temperature and wind speed readings connected to a timestamp. A source node with this stream subscribes to the URI 'http://...#AarhusWeatherData0'.

Table 2. CityBench: observation from AarhusWeatherData0 stream

Timestamp	hum	tempm	wspdm
2014-08-01T00:00:00	56.0	18.0	7.4

Listing 1.1 is the corresponding RDF mapping that stores the sensor and observation metadata of the weather data stream. It also contains bindings to fields from the underlying stream data e.g. 'AarhusWeatherData0.tempm'.

Listing 1.1. CityBench AarhusWeatherData0 RDF Mapping (abbreviated)

```
@prefix ssn:<http://purl.oclc.org/NET/ssnx/ssn#>
@prefix sao:<http://purl.oclc.org/NET/sao/>
@prefix ct:<http://.../citytraffic#>
@prefix ns:<http://../SampleEventService#>
@prefix iot:<http://iot.soton.ac.uk/s2s/s2sml#>
_:obs1 a ssn:Observation;
  ssn:observedProperty ns:Property-1;
  sao:hasValue"AarhusWeatherData0.tempm"^^iot:literalMap;
  ssn:observedBy ns:AarhusWeatherData0.
ns:Property-1 a ct:Temperature.
_:obs2 a ssn:Observation;
  ssn:observedProperty ns:Property-2;
  sao:hasValue "AarhusWeatherData0.hum"^^iot:literalMap;
  ssn:observedBy ns:AarhusWeatherData0.
ns:Property-2 a ct:Humidity.
_:obs3 a ssn:Observation;
  ssn:observedProperty ns:Property-3;
  sao:hasValue "AarhusWeatherData0.wspdm"^^iot:literalMap;
  ssn:observedBy ns:AarhusWeatherData0.
ns:Property-3 a ct:WindSpeed.
```

This mapping references various ontologies like the Semantic Sensor Network (SSN) Ontology[4] and City Traffic Ontology, providing a common way of describing sensor data, increasing interoperability. A formal definition of the mapping language is covered in our previous paper [16] and as a specification[5].

Similarly, a mapping is available for the stream of traffic data at various locations in the city. The traffic stream consists of fields like *avgSpeed* and *congestionLevel* connected to a *timestamp*.

Listing 1.2 shows Query 2 of CityBench expressed in the W3C recommended RSP-QL syntax for RSP engines[6], that at the time of writing no other engines support yet. When registered, it processes both traffic and weather streams for observations of traffic congestion level and weather (e.g. temperature), from a particular stretch of road, for the last 3 s. The semantic expressivity of the RDF graph query provides for interoperability, however, the underlying values *v1* to *v4* are actually from just 2 fields from each flattened stream.

Listing 1.2. CityBench Query 2: Finding the traffic congestion level and weather conditions of my planned journey (abbreviated)

```
SELECT ?v1 ?v2 ?v3 ?v4
FROM NAMED WINDOW :traffic ON <http://...#AarhusTrafficData158505> [RANGE PT3S]
FROM NAMED WINDOW :weather ON <http://...#AarhusWeatherData0> [RANGE PT3S]
WHERE {
  WINDOW :weather {
    ?obId1 a ssn:Observation;
     ssn:observedProperty ?p1;
     sao:hasValue ?v1;
     ssn:observedBy ns:AarhusWeatherData0.
    ?p1 a ct:Temperature. ... }
  WINDOW :traffic {
    ?obId4 a ssn:Observation;
     ssn:observedProperty ?p4;
     sao:hasValue ?v4;
     ssn:observedBy ns:AarhusTrafficData158505.
    ?p4 a ct:CongestionLevel. } }
```

4 http://www.w3.org/TR/vocab-ssn/.
5 https://github.com/eugenesiow/sparql2sql/wiki/S2SML.
6 https://www.w3.org/community/rsp/.

This query is then translated using the RDF mappings as explained in Sect. 4.1, distributed by Eywa's inverse-pub-sub, processed and returned to the client for processing as in Sect. 4.2 and are streamed to downstream applications. The entire architecture of Eywa's RSP engine is summarised in Sect. 4.3.

4.1 Query Translation

Firstly, the RSP-QL query is broken down into algebra as in Diagram 1.1.

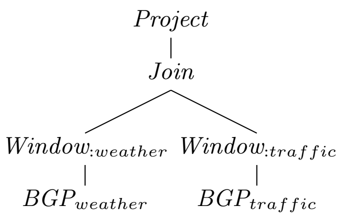

Diagram 1.1. Algebra of CityBench Query 2 from Listing 1.2

The algebra tree is traversed from leaf to root. At the $Window_{:weather}$ node, the basic graph pattern (BGP), $BGP_{weather}$, which compromises:

```
?obId1 a ssn:Observation;
  ssn:observedProperty ?p1;
  sao:hasValue ?v1;
  ssn:observedBy ns:AarhusWeatherData0.
?p1 a ct:Temperature. (...humidity and wind speed parts)
```

is matched against the *AarhusWeatherData0* RDF mapping from Listing 1.1. This matching can be done by any in-memory SPARQL engine and in our implementation we provide interfaces for popular open source engines[7]. The results of this step includes a result set bindings variable *?v1* from $BGP_{weather}$ to the field *AarhusWeatherData0.tempm* in the stream and so forth. Similarly, from $Window_{:traffic}$ we retrieve the binding of variable *?v4* to *AarhusTraffic-Data158505.congestionLevel* from the result set.

At the *Join* node, since no variables overlap from the *weather* and *traffic* windows, the result sets are passed upwards to the *Project* node.

At the *Project* node, variables $?v1$ to $?v4$ are projected. Any variable renaming/aliases will be taken care of in the projection as well. The simplified stream query expressed in Event Processing Language (EPL)[8] is shown in Listing 1.3:

[7] https://github.com/eugenesiow/sparql2sql/wiki/SWIBRE.
[8] http://www.espertech.com/products/esper.php.

Listing 1.3. CityBench Query 2 simplified and translated to EPL

```
SELECT AarhusWeatherData0.tempm AS v1 ,
  AarhusWeatherData0.hum AS v2 ,
  AarhusWeatherData0.wspdm AS v3 ,
  AarhusTrafficData158505.congestionLevel AS v4
FROM AarhusWeatherData0.win:time(3 sec) ,
  AarhusTrafficData158505.win:time(3 sec)
```

Query 5 which discovers the traffic congestion level on the road where a given cultural event is happening[9] has algebra (in Diagram 1.2):

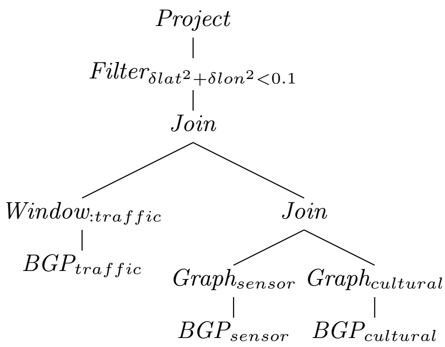

Diagram 1.2. Algebra of CityBench Query 5

$Window_{:traffic}$ produces a similar result set as in Query 2. However, $Graph_{sensor}$ and $Graph_{cultural}$ reference a static/graph source instead of a stream, hence the simplified translated syntax produced is SQL instead of EPL, though the method of executing against an RDF mapping is the same e.g. $Graph_{cultural}$ produces the following SQL when translated against the *AarhusCulturalEvents* mapping:

```
SELECT title, lat, lon FROM AarhusCulturalEvents
```

So at the first *Join* in the tree from the leaf, the SQL is integrated into the FROM clause of the translated EPL statements produced:

```
FROM sql:AarhusCulturalEvents [ ‘SELECT title, lat, lon
    FROM AarhusCulturalEvents‘ ] AS AarhusCulturalEsvents ,
sql:SensorRepository [ ‘SELECT lat, lon FROM SensorRepository‘ ] AS SensorRepository
```

At the next *Join* above, as there is a variable *?p2* present in both inbound nodes, a join is performed between the RDF metadata from $Window_{:traffic}$ and the previous *Join*, so the SQL within becomes (note the extra WHERE clause):

[9] https://github.com/eugenesiow/Benchmark/wiki/Q05.

```
... [ ‘SELECT lat, lon FROM SensorRepository
    WHERE propId=\’Property-b9f9...\’‘ ] AS SensorRepository,
AarhusTrafficData158505.win:time(3 sec)
```

Finally, a filter on the addition of the square delta of latitude and longitude constrains the traffic sensor and cultural event locations to within 0.1 units. All queries and corresponding translations can be found on our engine's wiki[10].

4.2 Query Distribution

Query distribution is the process whereby part of a query workload, ω_τ, is distributed from the client (where the query is registered and the results are expected), to the source node (where the data is produced or stored).

For example, for Query 2 in CityBench (Listing 1.2), at the *Project* operator at the top of the algebra tree, the engine tracks that *hum*, *tempm* and *wspdm* are the fields required from $Window_{:weather}$ while only *congestionLevel* is required from $Window_{:traffic}$. Hence, the projection of these fields is the work function, ω_τ, pushed to each source node producing the traffic and weather streams. The projection tree for the streams in Query 2 are shown below in Diagram 1.3.

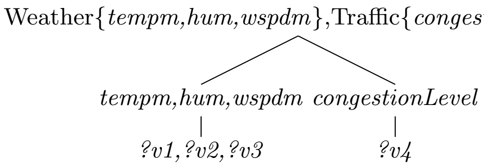

Diagram 1.3. Projection tree for Query 2 streams

Following RSP Query Translation (Sect. 4.1), the resulting simplified, translated EPL query is registered on the client. The root of the projection tree is distributed through Eywa's inverse-pub-sub mechanism to the subscribing source nodes of the IRIs of $Window_{:weather}$ and $Window_{:traffic}$. Hence, each source node pushes only an effective subset of each stream, γ_q, with only the projected fields to the client which further processes the EPL stream.

For static sources, the work function, ω_τ, also comprises SQL queries distributed to the source nodes. In Query 5 of CityBench, an SQL query:

```
SELECT lat , lon FROM SensorRepository WHERE propId=\’Property-b9f9...\’
```

is executed on the static *SensorRepository* source node, returning results to the client using a Java Database Connectivity (JDBC) connection.

10 https://github.com/eugenesiow/sparql2sql/wiki.

4.3 Eywa's RSP Architecture

Figure 2 shows the architecture of Eywa's RSP implementation consisting *client, source* and *broker* nodes. The *brokers* use lightweight ZeroMQ sockets and binary transmission protocol with minimal overhead to support publish-subscribe IRI topics. *Client* nodes receive queries in RSP-QL and projected streams via push from *source* nodes. They process complex events in queries using an EPL engine that registers translated queries with RDF mappings. *Source* nodes produce streams of time-series IoT data or store static data, e.g. sensor location data, cultural event data. A stream distributor component pushes projected streams, γ_q, to client nodes as required. Static queries utilise JDBC connections.

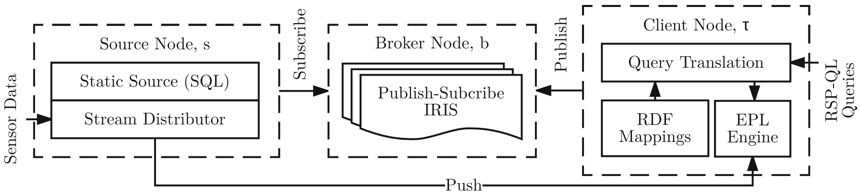

Fig. 2. Architecture of Eywa's RSP Fog computing framework

5 Experimentation Evaluating Eywa

The experiment uses the Smart City benchmark for Linked Data, CityBench [1], that features real-time datasets (e.g. vehicle traffic, parking, weather, pollution, etc.) from the city of Aarhus and streaming queries based on smart city application requirements (e.g. parking space finder, admin console). Linked Data goes towards fulfilling the interoperability metric and hence we evaluate Eywa's RSP framework by latency (performance) and scalability metrics.

The streaming and static source nodes, client nodes and broker nodes used resource-constrained lightweight computers, in the form of Raspberry Pi 2 Model B+s', suitable for inexpensive, mobile and broad Eywa Fog Computing networks. Each had 1 GB RAM, a 900 MHz quad-core ARM Cortex-A7 CPU and Class 10 SD Cards. 256 MB (the recommended and default 1/4th of system memory) was assigned to the Java Virtual Machine on Raspbian 4.1. Ethernet connections were used between the nodes for reliable transport.

We compared our approach, Eywa RSP, against state-of-the-art native RSP streaming engines CQELS [11] and C-SPARQL [2]. For each query, we varied the amount of concurrent queries and the number of data streams.

For each query and experimental configuration, we tested each engine for the latency and memory consumption. Tests were run for 15 min and averaged over 3 runs. The goal was to measure the performance (latency of queries) and scalability (the memory consumption while varying the number of data streams and concurrent queries). Additionally, to measure the benefits in scalability of our Fog Computing infrastructure, Eywa-RSP was tested and compared with client, source and broker on a single node against when they were across multiple nodes (3 nodes) with workload distribution.

6 Results and Discussion

6.1 Latency Evaluation

The latency of a query refers to the average time consumed by the engine between when a stream observation arrives and when the results from the query output are generated. Nodes are synchronised using the Network Time Protocol (NTP).

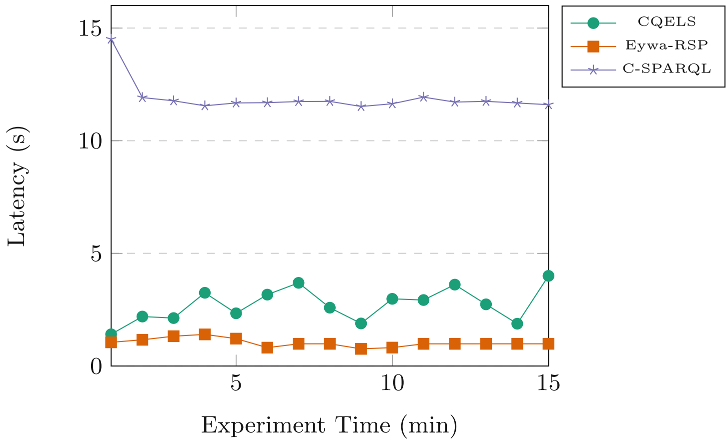

Fig. 3. Latency of CityBench Q1 across time

Figure 3 shows the latency over time of each of the engines for Query 1[11] which measures the traffic congestion level on two roads. As we can see, Eywa-RSP has the lowest latency and hence best query performance. All other queries show similar results, with Eywa having the lowest latency over time and are summarised by averaging over the 15 min interval as shown in Fig. 4.

The reason that Eywa has a lower latency than CQELS and C-SPARQL is because it abstracts metadata triples to mappings, which are processed in the query translation process as opposed to in the stream itself.

Diagram 1.4 shows the algebra of the translated query 1 from CityBench, registered on the client node. Π is the operator for retrieving the projected columns (e.g. congestionLevel) from each event in the stream window.

[11] https://github.com/eugenesiow/Benchmark/wiki/Q01.

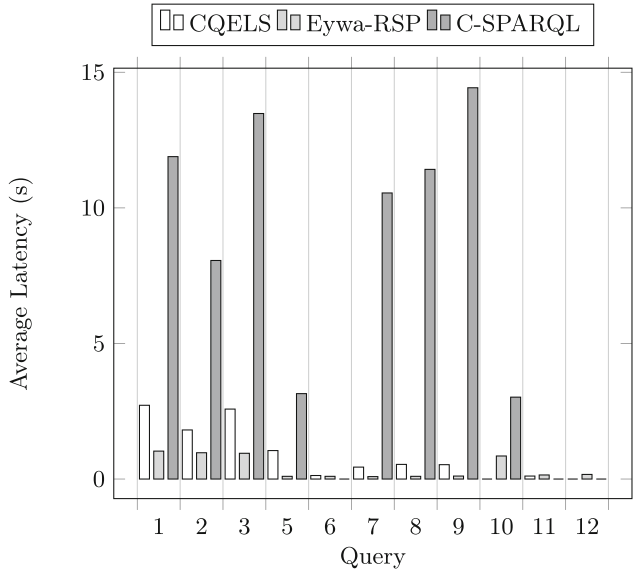

Fig. 4. Average latency of CityBench Queries on RSP Engines (Bars with zero values are queries that cannot be run on that engine (Eywa:4, CQELS:4,10,12, C-SPARQL:4,6,11,12))

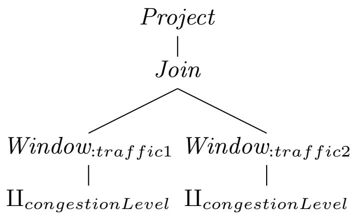

Diagram 1.4. Algebra of Query 1 for Eywa-RSP

The algebra for CQELS and C-SPARQL on the other hand, as shown in Diagram 1.5, requires the retrieval of, τ_{BGP}, for each event in the stream. The extra triples are shown in Listing 1.4, representing the observation metadata (lines 1 to 3) and relation to sensor metadata (line 4).

Listing 1.4. Additional metadata triples

```
?obId1 a ?ob.
?obId1 ssn:observedProperty ?p1.
?obId1 sao:hasValue ?v1.
?obId1 ssn:observedBy <...#AarhusTrafficData182955>.
```

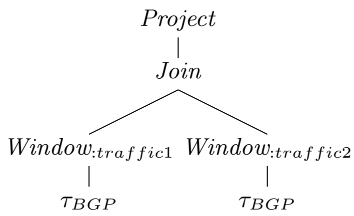

Diagram 1.5. Algebra of Query 1 for CQELS and C-SPARQL

It is more expensive in terms of overall latency when retrieving the extra triples and processing the query without workload distribution at the source. Hence, Eywa-RSP proves faster than the other tested engines.

6.2 Scalability Evaluation

Scalability evaluation looks at the amount of memory resources used by each RSP engine and the results when increasing the number of concurrent queries and increasing the number of data streams. The lower the memory resources used, the more scalable the system is, especially on resource-constrained nodes.

Figure 5 shows the memory consumption of each engine across time. Eywa-RSP streams a more concise format of just time-series data and only the necessary projected fields from the source node, as compared to the verbose set of triples, in Listing 1.5, for other engines. This is consistent for other queries in CityBench with the more complex Q5, a filter operation on 2 graphs and a stream, taking the most memory and Q11, checking for observations from weather sensors on a single stream, the least on all engines.

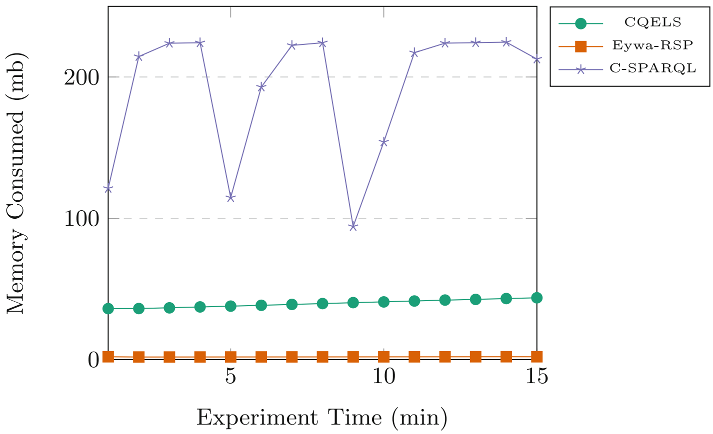

Fig. 5. Memory Consumed by CityBench Query 2 across time on client

Listing 1.5. Triples from a row/event of the Traffic Stream

```
:Property1 a ct:CongestionLevel.
:Observation1 a ssn:Observation;
  ssn:observedProperty :Property1;
  sao:hasValue "congestionLevel";
  ssn:observedBy ns:AarhusTrafficData158505;
  time:inXsdDateTime "timestamp".
:Observation2 a ssn:Observation;
  ssn:observedProperty :Property2;
  sao:hasValue "avgSpeed";
  ssn:observedBy ns:AarhusTrafficData158505;
  time:inXsdDateTime "timestamp".
```

As expected, when increasing the number of concurrent queries on each engine, Eywa-RSP once again achieved the lowest memory consumption. The memory consumed was also the only one consistent and did not increase over time like C-SPARQL and CQELS. This can be seen from Fig. 6 which shows the results of Query 5 at two different configurations of a single query and 20 concurrent queries. There were consistent results from these 2 configurations for other CityBench queries. Q2 and Q8 took up significantly more memory for C-SPARQL, Q1 and Q3 took up significantly more memory for CQELS while other queries took slightly more memory (about 10 MB). Eywa-RSP had a slight, stable 2–5 MB increase in memory consumption on increasing concurrent queries.

Figure 7 shows the memory consumption when increasing the number of pollution data streams from 2 to 5 in CityBench's query 10 which looks for the most polluted area in the city in real-time[12]. Eywa-RSP once again has the lowest memory consumption and C-SPARQL actually runs out of memory on the resource-constrained client in the 5 stream configuration just before 15 min.

Finally, we want to observe how much the Eywa fog computing framework, through the projection operator on the source node, actually improves

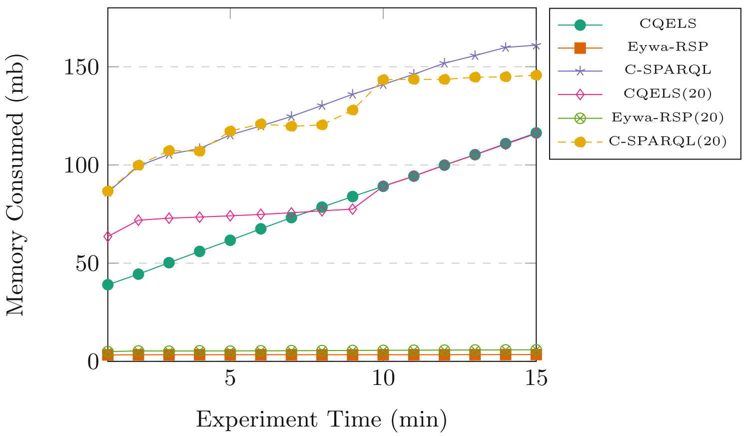

Fig. 6. Memory consumed when increasing # concurrent queries on CityBench Q5

[12] https://github.com/eugenesiow/Benchmark/wiki/Q10.

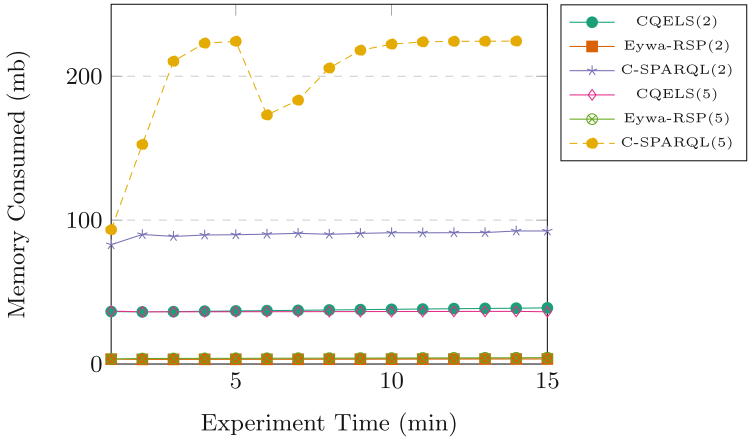

Fig. 7. Memory consumed by CityBench Q10 with varying data stream configurations

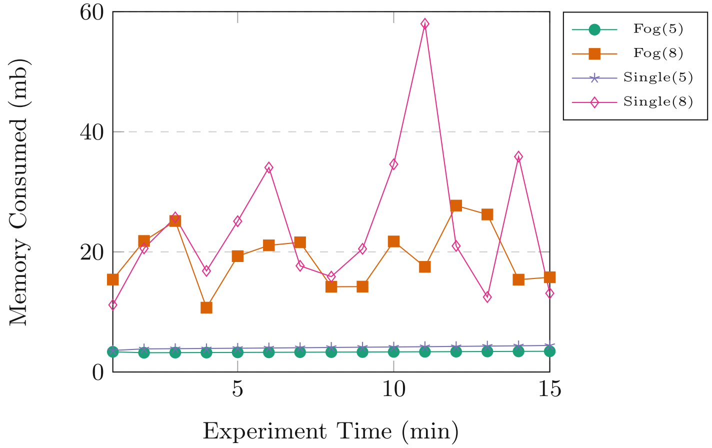

Fig. 8. Fog vs Single-Node: memory consumed while increasing streams on Q10

the memory consumption and scalability across an increasing amount of data streams. Figure 8 shows the amount of memory consumed by the distributed fog execution of the query and the single node execution of the query.

As the projection operator is distributed to the source nodes and only the applicable part of the stream is sent to the client, there is a significant difference in memory consumed when the number of incoming data streams are increased. At 2 and 5 streams for query 10, the memory consumed is similar for fog and single-node (non-fog) setups. When there were 8 streams on query 10, however, the fog computing approach consumed less memory. This is due to the flow of data being significantly large enough that the projection operation passed down to the source node of each stream also has a significant effect on the overall memory consumed. Hence, as the amount of streams increase, the Eywa's Fog Infrastructure provides greater scalability. Evaluation experiments forked from the CityBench benchmark harness are available on Github[13].

7 Fog Computing Within the Internet of the Future

The Internet has been a great motor of socio-economic activity in the past decade and the Internet of the Future has the potential to do even more. Like the fictional Eywa in Avatar, Eywa as a Fog Computing infrastructure has the potential to make the IoT and the planet more intelligent, more connected and at the same time more engaged and social while valuing privacy and security.

Learning from the web, that grew as a free, open and standards-based information space on top of the internet [4], the IoT needs to endeavour to support interoperability. Linked Data not only encourages interoperation through the use of URIs and common ontologies, it also represents concepts as machine-readable graphs [5] which goes towards solving challenges like discovery and data integration in the IoT. Fog Computing that serves as a layer between sensors and actuators and the cloud and possesses comparatively more compute and storage than sensors, serves as the entry level of interoperability and integration.

The current business model for many web services is that free content is exchanged for our personal data. However, as Berners-Lee states, 'as our data is held in proprietary silos, out of sight to us, we lose out on the benefits we could realise if we had direct control over this data and chose when and with whom to share it[14]'. Fog Computing on lightweight, distributed computers means that data that is collected from sensors and devices are stored and processed locally. As access controls evolve, specific privacy policies and access control with additional trust and fault tolerance mechanisms can be created [15].

Furthermore, as data is stored and processed by some applications locally, we minimise the need to shuttle data to and from the cloud, there is more quality of service guarantee for mission-critical IoT apps and there is less dependency on supporting high bandwidth global connections. In disaster management IoT scenarios, where last-mile connectivity is lost, having data locality and offline access is especially valuable. There are also performance benefits over storing encrypted data in traditional clouds as a means to maintain privacy because it is easier to perform processing (no need for crypto-processors or to apply special encryption functions) over data. When necessary and access controls permit, data can still be sent to the cloud for big data processing.

Finally, as the Internet and IoT advance, studying the social-technical aspects of the intersection between intelligent, cooperative autonomous machines in the IoT and human users is gaining importance. Fog Computing, as a distributed,

[13] https://github.com/eugenesiow/Benchmark.

[14] https://www.theguardian.com/technology/2017/mar/11/tim-berners-lee-web-inventor-save-internet.

interoperable application layer between the network, autonomous IoT end-devices, applications and human end-users can serve to (1) build and manage this social network, (2) facilitate information flow/sharing, (3) host applications and (4) serve as an observation platform for the study of these emerging networks.

Hence, we argue for the importance of research in Fog Computing technologies as part of the future, next-generation internet.

8 Related Work

Bonomi *et al.* [6] introduce Fog Computing as a platform for the Internet of Things, its defining principles, its utility for real-time streaming applications and the concept of *distributed orchestration*. These principles provide guidelines for the formal design and implementation of our framework which registers queries on streams and pushes results to real-time applications while introducing a novel inverse publish-subscribe model for workload management and orchestration.

In Wireless Sensor Networks literature, identifies base station nodes as "powerful devices with PC-like capabilities" which support the idea of Eywa Fog Computing nodes and are already widely-deployed. The innovation of a Fog Computing framework should come from managing the workload on and optimising performance of these nodes.

There has been previous work on RSP with the C-SPARQL [2] engine that supports continuous pull-based queries over RDF data streams by using Esper[15], a complex event processing engine, to form windows in which SPARQL queries can be executed on an in-memory RDF model. Another engine is CQELS [11], which is a purely native RSP, supporting both push and pull queries. Due to the 'white-box' approach, there is full control over query optimisation and execution. We compare against both of these engines.

Efficient SPARQL-to-SQL translation has been investigated by Rodriguez-Muro and Rezk [14] and Priyatna *et al.* [13]. The state-of-the-art reduces redundant self-joins and applies query containment and semantic query optimisation to translations. However, neither of the engines are designed to work on IoT streams or for Fog Computing scenarios yet.

There are a few benchmarks on streaming Linked Data including: SRBench [19], which in our previous work we have evaluated on [17], CityBench [1], which we compare against and LSBench [10] comprising social network stream data.

9 Conclusion and Future Work

The Internet of Things has huge research potential and Fog Computing, we argue, that is implementing a layer between the 'ground' and the 'cloud', can benefit from efficient, interoperable RDF stream processing (RSP) to support real-time applications on lightweight computers and networks. This does not replace, but seeks to complement large-scale processing and analytics in the

[15] http://www.espertech.com/products/esper.php.

cloud by providing collection, integration and continuous querying capabilities across distributed nodes. We go on to formally define and implement a Fog Computing infrastructure, Eywa, that utilises inverse-publish-subscribe and push as a distribution mechanism together with a query translation approach to RSP optimised for IoT time series data. In evaluation benchmarks, Eywa showed better performance and scalability as compared to state-of-the-art RSP approaches while preserving the interoperability of Linked Data.

Furthermore, with specialised distribution mechanisms for Fog Computing, we can maintain control, security and privacy on a per node level.

There is also potential to explore query distribution that pushes down a range of operators to the source nodes, Quality of Service guarantees and Service Level Agreements for source nodes and consensus algorithms for streams and sensors.

References

1. Ali, M.I., Gao, F., Mileo, A.: CityBench: a configurable benchmark to evaluate RSP engines using smart city datasets. In: Arenas, M., et al. (eds.) ISWC 2015. LNCS, vol. 9367, pp. 374–389. Springer, Cham (2015). doi:10.1007/978-3-319-25010-6_25
2. Barbieri, D.F., Braga, D., Ceri, S., Valle, E.D., Grossniklaus, M.: Querying RDF streams with C-SPARQL. ACM SIGMOD Record **39**(1), 20 (2010)
3. Barnaghi, P., Wang, W., Henson, C., Taylor, K.: Semantics for the internet of things: early progress and back to the future. Int. J. Semant. Web Inf. Syst. **8**(1), 1–21 (2012). http://dl.acm.org/citation.cfm?id=2441528
4. Berners-Lee, T., Fischetti, M.: Weaving the Web: The Original Design and Ultimate Destiny of the World Wide Web by its Inventor. Harper, San Francisco (2000). http://dl.acm.org/citation.cfm?id=554813
5. Bizer, C., Heath, T., Berners-Lee, T.: Linked data - the story so far. Int. J. Semant. Web Inf. Syst. **5**, 1–22 (2009). https://eprints.soton.ac.uk/271285/
6. Bonomi, F., Milito, R., Natarajan, P., Zhu, J.: Fog computing: a platform for internet of things and analytics. In: Bessis, N., Dobre, C. (eds.) Big Data and Internet of Things: A Roadmap for Smart Environments. SCI, vol. 546, pp. 169–186. Springer, Cham (2014). doi:10.1007/978-3-319-05029-4_7
7. Bonomi, F., Milito, R., Zhu, J., Addepalli, S.: Fog computing and its role in the internet of things. In: Proceedings of the 1st Workshop on Mobile Cloud Computing, pp. 13–16 (2012). http://dl.acm.org/citation.cfm?id=2342513
8. Buil-Aranda, C., Hogan, A., Umbrich, J., Vandenbussche, P.-Y.: SPARQL web-querying infrastructure: ready for action? In: Alani, H., et al. (eds.) ISWC 2013. LNCS, vol. 8219, pp. 277–293. Springer, Heidelberg (2013). doi:10.1007/978-3-642-41338-4_18
9. Compton, M., Barnaghi, P., Bermudez, L., García-Castro, R., Corcho, O., Cox, S., Graybeal, J., Hauswirth, M., Henson, C., Herzog, A., Huang, V., Janowicz, K., Kelsey, W.D., Le-Phuoc, D., Lefort, L., Leggieri, M., Neuhaus, H., Nikolov, A., Page, K., Passant, A., Sheth, A., Taylor, K.: The SSN ontology of the W3C semantic sensor network incubator group. J. Web Semant. **17**, 25–32 (2012). http://linkinghub.elsevier.com/retrieve/pii/S1570826812000571
10. Le-Phuoc, D., Dao-Tran, M., Pham, M.-D., Boncz, P., Eiter, T., Fink, M.: Linked stream data processing engines: facts and figures. In: Cudré-Mauroux, P., et al. (eds.) ISWC 2012. LNCS, vol. 7650, pp. 300–312. Springer, Heidelberg (2012). doi:10.1007/978-3-642-35173-0_20

11. Le-Phuoc, D., Dao-Tran, M., Xavier Parreira, J., Hauswirth, M.: A native and adaptive approach for unified processing of linked streams and linked data. In: Aroyo, L., Welty, C., Alani, H., Taylor, J., Bernstein, A., Kagal, L., Noy, N., Blomqvist, E. (eds.) ISWC 2011. LNCS, vol. 7031, pp. 370–388. Springer, Heidelberg (2011). doi:10.1007/978-3-642-25073-6_24
12. Mell, P., Grance, T.: The NIST definition of cloud computing. Technical report (2011). http://csrc.nist.gov/publications/nistpubs/800-145/Sp800-145.pdf
13. Priyatna, F., Corcho, O., Sequeda, J.: Formalisation and experiences of R2RML-based SPARQL to SQL query translation using morph. In: Proceedings of the 23rd International Conference on World Wide Web, pp. 479–489 (2014). http://dl.acm.org/citation.cfm?id=2567981
14. Rodriguez-Muro, M., Rezk, M.: Efficient SPARQL-to-SQL with R2RML mappings. Web Semant.: Sci. Serv. Agents World Wide Web **33**, 141–169 (2015). doi:10.1016/j.websem.2015.03.001
15. Roman, R., Zhou, J., Lopez, J.: On the features and challenges of security and privacy in distributed internet of things. Comput. Netw. **57**(10), 2266–2279 (2013). http://www.sciencedirect.com/science/article/pii/S1389128613000054
16. Siow, E., Tiropanis, T., Hall, W.: Interoperable and efficient: linked data for the internet of things. In: Bagnoli, F., et al. (eds.) INSCI 2016. LNCS, vol. 9934, pp. 161–175. Springer, Cham (2016). doi:10.1007/978-3-319-45982-0_15
17. Siow, E., Tiropanis, T., Hall, W.: SPARQL-to-SQL on internet of things databases and streams. In: Groth, P., Simperl, E., Gray, A., Sabou, M., Krötzsch, M., Lecue, F., Flöck, F., Gil, Y. (eds.) ISWC 2016. LNCS, vol. 9981, pp. 515–531. Springer, Cham (2016). doi:10.1007/978-3-319-46523-4_31
18. Stankovic, J.A.: Research directions for the internet of things. IEEE Internet Things J. **1**(1), 3–9 (2014). http://ieeexplore.ieee.org/document/6774858/
19. Zhang, Y., Duc, P.M., Corcho, O., Calbimonte, J.-P.: SRBench: a streaming RDF/SPARQL Benchmark. In: Cudré-Mauroux, P., et al. (eds.) ISWC 2012. LNCS, vol. 7649, pp. 641–657. Springer, Heidelberg (2012). doi:10.1007/978-3-642-35176-1_40

Large-Scale Open Corporate Data Collection and Analysis as an Enabler of Corporate Social Responsibility Research

Vasiliki Gkatziaki(✉), Symeon Papadopoulos, Sotiris Diplaris, and Ioannis Kompatsiaris

Information Technologies Institute (ITI), CERTH -ITI, Thessaloniki, Greece
vasgat@iti.gr

Abstract. During the last years, citizens and transparency initiatives put increasing pressure on governments, organizations, and companies to be more transparent and to publicize information pertaining to their operations. Although several organizations have started engaging in open data practices, data quality, structure and availability is still highly inconsistent across organizations, which makes it challenging and effort-intensive to obtain and analyze large-scale high-quality datasets. To this end, this paper examines how publicly available financial and corporate data can be leveraged to extract useful inferences regarding the financial and social performance of companies. Numerous reports have been collected from the Securities Exchange Commission (SEC) and analyzed to study hypotheses regarding the corporate practices and social responsibility of companies.

Keywords: Open data · Information retrieval · Corporate Social Performance · eXtensible Business Reporting Language (XBRL)

1 Introduction

During the last decade, important efforts have been expended towards ensuring transparent reporting regarding the Environmental, Social and Governance performance of companies. Governments, shareholders and other stakeholders require companies to become more transparent regarding their financial, social and environmental activities. For instance, the Securities and Exchange Commission (SEC), a government commission created by the US Congress with the goal of protecting investors from deceitful and manipulative market practices, requires US public companies to file statements regarding their financial performance in order to protect investors. In 2008, SEC adopted a new rule [21] that enforces public companies to disclose their financial statements using the eXtensible Business Reporting Language (XBRL) format[1] in an attempt to increase financial transparency.

[1] XBRL (https://www.xbrl.org/) is a global standard for business reporting that allows the digitization of financial statements.

I. Kompatsiaris et al. (Eds.): INSCI 2017, LNCS 10673, pp. 266–280, 2017.
https://doi.org/10.1007/978-3-319-70284-1_21

WikiRate[2] is a platform for collecting and analyzing information about companies' Environmental, Social and Governance (ESG) performance and it aims to make corporations more transparent, reactive and ethical by making data about their ESG performance available to everyone [17]. Numerous metric points, coming from different sources, have been collected and integrated into the WikiRate platform. Moreover, a plethora of financial metric points have been extracted from 10-K filings[3] and integrated into the platform.

In this paper, we investigate how open financial data can be leveraged to extract useful inferences regarding the social performance of companies. We collected 39,029 10-K reports in XBRL format from SEC and after processing these reports, we managed to extract more than 529,000 financial facts related to different aspects of the financial performance of companies in different reporting years. Additionally, we collected about 3,442 Conflict Minerals reports. Then, we performed data analysis on two datasets to study research hypotheses regarding specific aspects of Corporate Social Responsibility (CSR). The first dataset, we studied, comprised 25,500 observations regarding 7,700 companies, while the second comprised 465 observations regarding 465 companies and contained financial facts extracted from XBRL reports as well as green scores defined by Newsweek in 2016[4].

The main contributions of this work can be summarized as follows:

– *Collecting a large amount of data from the Web in relation to the financial performance of companies*: A large database of over 500,000 financial facts about more than 50,000 companies was extracted from 10-K filings (available by SEC). Additionally, a REST API was developed to make the collected data available to third parties for further research.
– *Demonstrating the value of open data for CSR research:* We argue that financial open data can lead to useful inferences regarding the social performance of companies. We support our claim by testing four hypotheses on two datasets that integrate different types of data regarding the social and environmental performance of companies.

Figure 1 presents an overview of the methodology adopted in this paper. This comprises three key steps: *Data Collection*, *Data Integration* and *Data Analysis*. *Data Collection* is responsible for constructing appropriate wrappers for the given Web sources, executing the constructed wrappers and handling the extracted data. *Data Integration* results in the construction of appropriate datasets by integrating data coming from different sources. Finally, *Data Analysis* on the created datasets performs *descriptive statistics* and *statistical inference* to verify selected research hypotheses.

[2] http://wikirate.org/.

[3] A 10-K is a comprehensive summary report of a company's performance that must be submitted annually to the SEC.

[4] Newsweek, an American news magazine, publishes yearly rankings for the 500 largest publicly-traded US companies based on their overall environmental performance. For 2016, those are available on http://www.newsweek.com/green-2016.

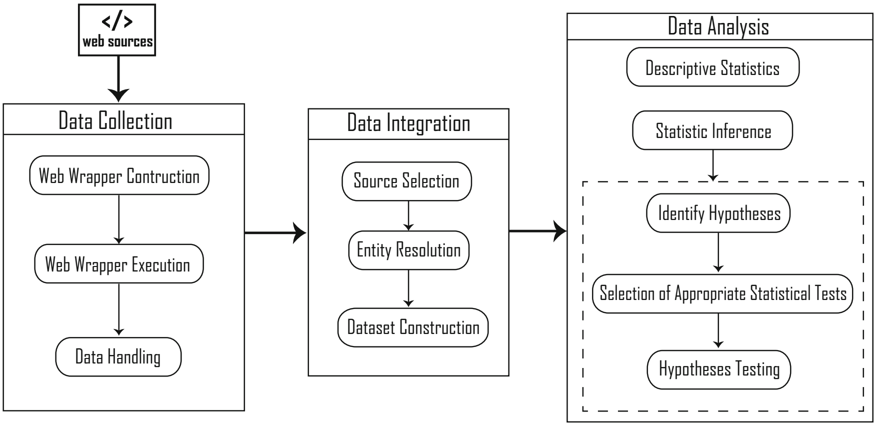

Fig. 1. Research methodology.

2 Related Work

The concept of Corporate Social Responsibility (CSR) emerged in the United States almost 60 years ago [2]. Over the years, numerous researchers tried to provide different definitions of CSR, to achieve better understanding and to explore different dimensions of the concept [5,7,11,18,22]. For instance, Wanderley et al. [24] investigated 127 corporations and how their CSR communication on the Web is influenced by the country or industry. They concluded that CSR disclosure on the Web is significantly influenced by the country where the corporation operates but less by the industry sector. Cai et al. [4] researched 475 companies from controversial industries in terms of producing dangerous products for humans, environment or society to study the importance of CSR engagement. Their results indicated that CSR engagement even for those industries is important and helped them increase their value. Bauman and Skitka [1] studied how CSR affects employees and developed a framework that identifies the correlation between employees' needs and companies' actions. Parguel et al. [20] addressed the question of how sustainable ratings affect consumers and discovered that such ratings significantly influence consumers to make more effective assessments on companies. Namkung and Jang [19] conducted a survey to discover which characteristics play an important role on consumers to pay more for green practices in restaurants. Finally, Loughran et al. [13] studied the use of ethics-related terms in 10-K reports obtained from SEC and detected an inconsistency between companies' social performance and the use of such terms.

The relationship between Corporate Social Performance (CSP) and Corporate Financial Performance (CFP) has received considerable attention over the years. Numerous studies have been performed to investigate this relationship. Bragton and Marlin [3] were among the first to study the connection between CSP and CFP. They investigated the assumption that more strict pollution controls translate to less profit. Their outcomes negate the hypothesis that environmental virtue and profitability are incompatible goals. Bromiley and Marcus [14] studied the assumption that stock market reaction to problematic behaviors can function as an instrument of social control to discourage such behaviors. Their findings indicated that stock market as a medium of social control is inadequate. Waddock and Graves [23] tried to use a more representative measurement approach for CSP by taking into consideration different aspects and discovered a positive connection between CSP and CFP. Margolis et al. [15] performed a meta-analysis on 167 studies to define which CSR aspects are more effective over the financial performance of a firm and they discovered that there is a small positive effect of CSP over CFP. Cheng et al. [6] researched how CSP affects a firm's ability to access capital. They concluded that companies with superior social performance have lower capital limitations. Eckbo et al. [9] studied the effect of gender-balancing in the board of directors in the case of Norwegian companies. The results showed that gender-quota has neutral effect on a firm's value.

With the emergence of Web 2.0 and the prevalence of the Web as a global communication medium, governments are increasingly demanding companies to disclose data regarding their financial and social performance in order to protect citizens and initiate an open dialogue. As a result, large amounts of information about companies' financial and social performance are available online by governments, NGOs and other initiatives. In this paper, we study how open financial data disclosed by US companies to the SEC can lead to useful inferences regarding specific aspects of their social performance. Additionally, an integrated dataset is studied that combines data from two different sources and investigates the relationship of CSP and CFP.

3 Data Collection, Integration and Publishing

3.1 Company Data Model and Metrics

Figure 2 illustrates the developed data model, comprising the following entities:

- `companies`: US companies, each associated with a name, a Central Index Key[5] (CIK) and a Standard Industrial Classification (SIC) code.
- `metrics`: pieces of information related to companies; each metric is defined by its name, value, source, the year to which it refers and the CIK referring to the associated company.

The underlying database, mongoDB, is a schema-free document database, which offers the required flexibility allowing to add more fields based on the needs of each extracted metric. The serialization format is BSON (binary JSON), which offers higher storage efficiency.

[5] https://www.sec.gov/edgar/searchedgar/cik.htm.

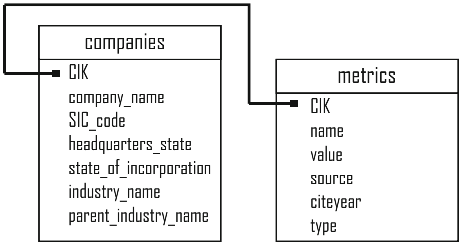

Fig. 2. Data model used for storing the collected data.

3.2 Data Collection and Extraction from SEC

This study has been based on data collected from the Securities and Exchange Commission (SEC). SEC supports transparency and requires public companies to file registration statements, periodic reports and other securities forms electronically. Most of these high-value data are received, processed and disseminated through the Electronic Data Gathering, Analysis and Retrieval (EDGAR) system. We obtained from the public database of EDGAR two types of report:

– *Conflict Minerals Report:* companies are obliged to disclose in these documents their use of conflict minerals considered necessary to production.
– *Form 10-K:* these documents make available audited financial statements regarding the annual financial performance of the company.

Collecting the aforementioned reports in large scale was a challenging task, which necessitated the development of a custom Web wrapper, i.e. a highly automated piece of software that iterates through a number of Web pages of interest and extracts structured information from each page [10]. Algorithm 1 specifies the operation of the developed SEC Wrapper. Given as input the set of all available Standard Industry Codes (SIC) and the report types of interest, the wrapper searches for companies, and for each company it extracts relevant information, such as the company name, the Central Index Code, the state of incorporation, etc. from the appropriately selected DOM[6] elements. In addition to basic company information, all reports of the selected type (e.g., Form 10-K and Conflict Minerals) are extracted from the respective company page. To specify the target DOM elements, i.e. the placeholders of interesting information within the fetched HTML pages, we composed extraction rules in the form of *CSS selectors*. The implementation of the wrapper was based on the JSoup[7] Java library. Using the wrapper, we collected data for 50,219 US companies, 3,442 Conflict Minerals Reports and 162,349 Form 10-K filings.

SEC is also supporting efforts to make data easy to process in a programmatic way. Numerous reports in EDGAR are available in this format. Due to

[6] The Document Object Model (DOM) is a programming interface for HTML and XML documents.
[7] http://jsoup.org/.

Algorithm 1. SEC Wrapper

Input: set of standard industry codes S_{SIC}; report types S_{RT}

foreach $sic \in S_{SIC}$ **do**
    companies $\leftarrow$ lookup(sic)
    foreach *company* $c \in$ *companies* **do**
        $\mathrm{url}_c \leftarrow$ create_company_url(c)
        extract_company_info(url_c)
            (i) extract content from defined DOM elements
            (ii) post-processing of the extract content if needed
            (iii) store company data
        foreach $r \in S_{RT}$ **do**
            extract_reports(url_c, r)
                (i) extract content from defined DOM elements
                (ii) store report

SEC, companies have to use the Generally Accepted Accounting Principles (usgaap) for financial reporting, i.e. a standard framework of guidelines for financial accounting, including standards, conventions and rules to be followed by accountants.

XBRL supports the representation of complex financial statements. In total, we managed to collect 39,029 Form 10-K filings in XBRL format. For extracting financial facts related to a company from the collected reports, we leveraged the DOM tree of the instance document of the Form 10-K. Overall, we managed to extract a plethora of financial data points related to 27 metrics (Table 1). Out of those, our data analysis focused on the four metrics of Table 2.

Table 1. Metrics extracted from XBRL reports.

Metrics
Advertising Expense, Assets, Comprehensive Income Net of Tax, Cost of Goods Sold, Cost of Goods and Services Sold, Current Foreign Tax Expense Benefit, Current Income Tax Expense Benefit, Deferred Income Tax Expense Benefit, Deferred Tax Liabilities Undistributed Foreign Earnings, Effective Income Tax Rate Continuing Operations, Effective Income Tax Rate Reconciliation at Federal Statutory Income Tax Rate, Good Will, Gross Profit, Income Loss From Continuing Operations Before Income Taxes Foreign, Income Tax Expense Benefit, Income before income taxes, Profit, Net Income Attributable to Non Controlling Interest, Net Income Attributable to the company, Research and Development Expense, Revenues, Sales Revenue Goods Net, Sales Revenue Net, Selling and Marketing Expense, Undistributed Earnings Of Foreign Subsidiaries, Unrecognized Tax Benefits Reductions Resulting from Lapse of Applicable statute of Limitations, Unrecognized Tax Benefits That would impact Effective Tax Rate

Table 2. Metrics extracted from XBRL reports and used in our data analysis.

Metric	Details
Profits	The consolidated profit or loss of the company for the reported period
Good Will	Good Will is the value of a company's brand name and it depends on the customer base, customer relations, employee relations, its patents and technology
Research & Development (R&D) Expense	The costs related to R&D for the reporting period which aims on discovering new knowledge with the ambition to exploit such knowledge on developing new products or services
Undistributed earnings of foreign subsidiaries	Undistributed earnings of foreign subsidiaries that are intended to be indefinitely reinvested for the reported period

3.3 Data Availability

We developed a RESTful API on top of the collected data in order to make them available to third parties for further study. The API, which is available on the endpoint http://easie.iti.gr/sec_dataset/, returns data in JavaScript Object Notation (JSON) and offers two basic methods are available:

- *GET companies*[8], which is responsible for returning a collection of companies by querying their names with a search term. If no query is defined then the method returns all available companies.
- *GET metrics*[9], which is responsible for returning a collection of metrics data points by querying their name with a search term. If no query is defined then the method returns all available metrics.

There are several parameters that users can specify to limit the results on both methods. These are specified in Tables 3 and 4.

Table 3. Available options of *GET companies* method.

Parameter	Description
`q`	Query companies given a search term
`CIK`	Search for a company given its central index key
`SIC`	Search for companies based on industry code
`headquarters_state`	Search for companies based on their headquarters location
`state_of_inc`	Search for companies based on their state of incorporation
`page`	Pagination of the results (100 results per page)

[8] http://easie.iti.gr/sec_dataset/companies.
[9] http://easie.iti.gr/sec_dataset/metrics.

Table 4. Available options of *GET metrics* method.

Parameter	Description
`q`	Query metrics given a search term
`CIK`	Search for metrics related to a company based on its central index key
`citeyear`	Search metrics related to a specific year
`page`	Pagination of the results (100 results per page)

4 Data Analysis

We carried out our data analysis on two datasets. We first provide some descriptive statistics regarding our primary dataset obtained from SEC and then perform statistical inference on both datasets to study selected research hypotheses regarding the CSR performance of companies.

4.1 Descriptive Statistics

The dataset obtained from SEC comprises 50,219 US companies. A large number of the collected companies (43%) belong to Finance, Insurance and Real Estate Industry, 20% to Manufacturing, 12% to Services and 7% to Transportation, Communications, Electric, Gas and Sanitary Service. It is common practice among companies to have a different location for their headquarters compared to the location of incorporation. Even though most of the collected companies are located in California (16%), New York (16%) and Texas (8%), a large percentage of these companies are incorporated in Delaware (52%) and Nevada (22%), most likely due to their highly attractive tax rates. Delaware is considered as a domestic tax haven and its paradigm has been extensively studied [8,25].

We also collected about 3,500 Conflict Minerals Reports related to 1,265 companies. The main industry the companies report in those is Manufacturing (76%), followed by Retail Trade (7%) and Services (7%) (Fig. 3). Figure 4 depicts

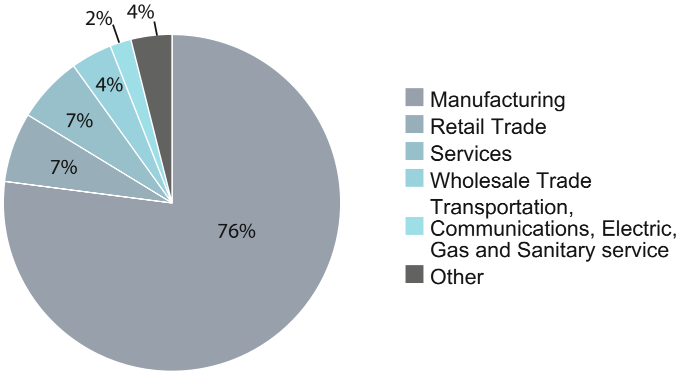

Fig. 3. Companies industry reporting conflict minerals.

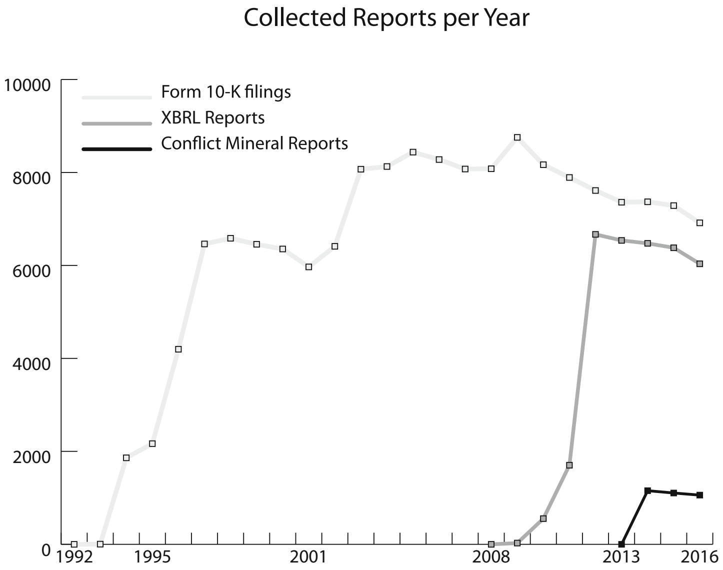

Fig. 4. Overview of the different types of collected reports per year.

the collected reports per year. It is noteworthy that companies started publishing financial annual reports in 1993 and over the years an increasing number of companies publish financial statements. In 2009, 27 companies started using XBRL in Form 10-K and 10-Q reporting. Companies started using XBRL more extensively in 2013 (6,668 companies published Form 10-K in XBRL format).

4.2 Statistical Inferences on CSR Performance

In this section, we are going to demonstrate the usefulness of the available financial open data for inference extraction regarding the social performance of companies. In particular, we perform data analysis on two datasets. The first contains more than 25,500 financial facts (from XBRL reports) for about 7,700 companies, and the second combines data about the environmental performance of 465 companies based on their green score in Newsweek 2016 Green List.

Financial facts dataset. Over the years, large corporations developed tax avoidance strategies to reduce their effective tax rate. These strategies are legitimate and companies adopting such strategies are not associated with inappropriate actions. For instance, the US government is actually incentivizing companies and individuals to increase their spending on charities by providing tax deductions in relation to such spending. In the following, we are going to study a different aspect of tax avoidance. Undistributed earnings of foreign subsidiaries (also known as permanently reinvested earnings) have been increasing over the

Table 5. Non-parametric correlation of the undistributed earnings to foreign subsidiaries to profits and R&D expenses. Each variable is perfectly correlated with itself (r = 1) and to this end the significance value is not calculated and is annotated in the table with ".".

		v1	*v2*	*v3*
Undistributed earnings of foreign subsidiaries (*v1*)	Correlation coefficient	1.000	0.637**	0.582**
	Sig. (2-tailed)	.	0.000	0.000
	N	3643	3512	2039
Profit (*v2*)	Correlation coefficient	0.637**	1.000	0.135**
	Sig. (2-tailed)	0.000	.	0.000
	N	3512	24278	7786
Research and development expense (*v3*)	Correlation coefficient	0.582**	0.135**	1.000
	Sig. (2-tailed)	0.000	0.000	.
	N	2039	7786	8040

**Correlation is significant at the 0.01 level (2-tailed).

last few years[10]. Corporations re-invest these earnings to their subsidiaries to avoid taxation. These practices can eventually lead to economic contraction and a shift of jobs overseas [16].

Laplante and Nesbitt [12] investigated the relationship between the trapped cash in foreign subsidiaries, the permanently reinvested earnings and the foreign cash. They found that the three terms are intersecting and also discovered that companies which have trapped cash are more likely to invest heavily in Research & Development. To this end, we are going to study the relationship of the undistributed earnings of foreign subsidiaries with not only profits but also with R&D Expenses. In this context, we are studying the following two hypotheses:

- **Hypothesis 1.** There is no correlation between the reported undistributed earnings of foreign subsidiaries, the profits and the R&D expenses.
- **Hypothesis 2.** There is no difference between the means of reported undistributed earnings of foreign subsidiaries between the years 2013, 2014, 2015 and 2016.

To study the first hypothesis, we used correlation tests and more specifically the Spearman's correlation coefficient that is appropriate for non-normal data. Both Kolmogorov-Smirnov and Shapiro-Wilk tests for normality showed that the undistributed earnings of foreign subsidiaries, profits and R&D expenses are significantly non-normal with $p < .001$. The Spearman's rho correlation coefficient showed a significant positive association between all three variables and resulted in the rejection of the null Hypothesis 1. The corresponding correlation coefficients are presented in Table 5. It is noteworthy that the relationship

[10] https://blogs.wsj.com/cfo/2015/04/24/companies-leaving-trillions-in-cash-overseas/.

Table 6. Top 10 companies ranked by the amount of profits in 2016. UEFS stands for Undistributed Earnings of Foreign Subsidiaries, and FIRE stands for Finance, Insurance and Real Estate.

Company	Industry	State of Inc.	Profits (in $M)	UEFS (in $M)
Apple Inc.	Manufacturing	California	$ 45,687	$ 109,800
JP Morgan Chase & Co.	FIRE	Delaware	$ 24,733	$ 38,400
Berkshire Hathaway Inc.	FIRE	Delaware	$ 24,427	$ 12,400
Wells Fargo & Co.	FIRE	Delaware	$ 22,045	$ 2,400
Alphabet Inc.	Services	Delaware	$ 19,478	$ 60,700
Bank of America Corp.	FIRE	Delaware	$ 17,906	$ 17,800
Microsoft Corp.	Services	Washington	$ 16,798	$ 124,000
Wal Mart Stores Inc.	Retail Trade	Delaware	$ 15,080	$ 26,100
Citigroup Inc.	FIRE	Delaware	$ 14,975	$ 47,000
Gilead Sciences Inc.	Manufacturing	Delaware	$ 13,488	$ 37,600

of R&D expenses appears to be stronger to undistributed earnings of foreign subsidiaries than to profits in terms of the coefficient value. This indicates that companies that permanently reinvest earnings of foreign subsidiaries are likely to invest heavily on R&D. Strong positive association between the undistributed earnings of foreign subsidiaries and company profits are also observed. Hence, we could claim that companies with high profits tend to adopt tax avoidance practices in terms of higher undistributed earnings of foreign subsidiaries.

Table 6 presents the top 10 companies in terms of profits in 2016 along with their undistributed earnings of foreign subsidiaries. The highest profits were reported by Apple Inc. followed by JP Morgan Chase & Co. We notice that in many cases the undistributed earnings of foreign subsidiaries exceed the amount of profits. The amount of the undistributed earnings of foreign subsidiaries for Apple Inc. was almost 2.5 times greater than the reported profits while for Microsoft was 7 times.

Continuing, the one-way analysis of variance (one-way ANOVA) test is commonly used to research the existence of statistically significant differences among the means of two or more groups and it would be a good candidate for researching the second hypothesis. However, our data failed the assumptions of normality for each category of the independent variable (year) as well as the homogeneity of variances assumption. To test normality within the groups (group by year), we used Kolmogorov-Smirnov's test and the undistributed earnings of foreign subsidiaries in reporting years 2013 ($D(859) = .401, p < .001$), 2014 ($D(911) = .407, p < .001$), 2015 ($D(939) = .403, p < .001$) and 2016 ($D(934) = .414, p < .001$) were found significantly different from the normal distribution. To test variance homogeneity, we used Lavene's test and the variances of the undistributed earnings of foreign subsidiaries in the four groups (grouped by year) were found significantly different, $F(3, 3639) = 5.32, p < .01$ and thus the homogeneity of variance could not be assumed.

In order to detect differences between the groups we used the non-parametric test of Kruskal-Wallis, equivalent to the one-way ANOVA. The results showed that the undistributed earnings of foreign subsidiaries were significantly different over the years, $H(3) = 14.867, p < .01$ and that led to the rejection of the null hypothesis 2. Mann-Whitney tests were used to reveal between which years there was a significant difference in the undistributed earnings of foreign subsidiaries. A Bonferroni correction was applied and thus all effects are reported at a .0083 level of significance. Significant differences in the undistributed earnings of foreign subsidiaries were detected only regarding years 2013 to 2015 ($U = 372062.5, r = -.067$) and 2016 ($U = 361298.0, r = -.086$). We can conclude that the undistributed earnings of foreign subsidiaries in 2015 and 2016 were significantly higher compared to 2013. Finally, Jonckheere's test revealed a significant trend in the data: the median of the undistributed earnings of foreign subsidiaries are significantly increased over the years ($J = 2621883.0, z = 3.790, r = .063$).

To sum up, the above analysis provides evidence that over the years profitable companies try to find ways to avoid taxation.

Newsweek green rankings dataset. Newsweek publishes rankings based on the environmental performance of the 500 largest publicly-traded Global and US companies since 2009. They score companies based on their performance on eight specific environmental aspects. In this section, we are going to study if and how we can extract useful inferences by combining data residing at different sources. More specifically, we combined the 2016 Newsweek green scores of 465 US companies with several financial statements extracted from 10-K filings. We are going to study the following two hypotheses:

- **Hypothesis 3.** There is no correlation between the Newsweek green score with Profits, Good Will and R&D Expenses.
- **Hypothesis 4.** There is no difference in the environmental performance (in terms of Newsweek green score) of companies reporting conflict minerals with those that do not.

To study Hypothesis 3, we used correlation tests for discovering relations between the four variables of interest (Newsweek green score, Profits, Good Will and R&D Expenses). The selection of appropriate correlation tests required to check for normality inside our sample with Shapiro-Wilk's test of normality. The distributions of all four variables were found to be significantly non-normal with $p < .001$ in all cases. Thus, we select Spearman's correlation coefficient to discover correlations and the strength of the existing relationships between them. Spearman's correlation coefficient in Table 7 indicates that there is strong positive correlation between all of them. This suggests that the higher the Newsweek green score is the higher the Profits, the Good Will and the R&D Expenses. Note that the highest coefficient occurs in the case of Profits with $r = .401$. Consequently, we could claim that big companies, in terms of financial performance, tend to adopt more sustainable practices.

Table 7. Non-parametric correlation of the Newsweek green score to Profits, Good Will and R&D Expenses. Each variable is perfectly correlated with itself (r = 1) and to this end the significance value is not calculated and is annotated in the table with ".".

		v1	*v2*	*v3*	*v4*
Newsweek green score (*v1*)	Correlation coefficient	1.000	0.401**	0.197**	0.232**
	Sig. (2-tailed)	.	0.000	0.000	0.005
	N	465	416	383	145
Profit (*v2*)	Correlation coefficient	0.401**	1.000	0.478**	0.437**
	Sig. (2-tailed)	0.000	.	0.000	0.000
	N	416	416	371	141
Good Will (*v3*)	Correlation coefficient	0.197**	0.478	1000	0.459**
	Sig. (2-tailed)	0.000	0.000	.	0.000
	N	383	371	383	136
Research and development expense (*v4*)	Correlation coefficient	0.232**	0.437**	0.459**	1.000
	Sig. (2-tailed)	0.005	0.000	0.000	.
	N	145	141	136	145

**Correlation is significant at the 0.01 level (2-tailed).

To test Hypothesis 4, we performed Mann-Whitney's non-parametric test since our data violated the normality assumption. Results indicate that the environmental performance of companies that report conflict minerals is significantly higher than those companies that do not ($U = 20358, r = -.15, p < .01$).

5 Conclusions

Over the years, more data are available online regarding the financial and social performance of companies. In this work, we examined how financial open data can be harnessed to extract useful inferences regarding specific aspects of Corporate Social Responsibility. We managed to collect numerous financial and conflict minerals reports from the SEC database and to extract a plethora of financial facts regarding the financial performance of companies. We performed data analysis on two datasets and tested four hypotheses regarding corporate tax avoidance and environmental performance. Integrating data from different sources is a challenging task but it could lead to important insights. To this end, we make the collected data available to third parties though a Web API for further use and investigation.

Future work could focus on researching more complex hypotheses as well as on integrating data from different sources related not only to the financial but also to the social performance of companies.

Acknowledgements. This work has been supported by the WikiRate and ChainReact projects, partially funded by the European Commission under contract numbers FP7-609897 and H2020-687967 respectively.

References

1. Bauman, C.W., Skitka, L.J.: Corporate social responsibility as a source of employee satisfaction. Res. Organ. Behav. **32**, 63–86 (2012)
2. Bowen, H.R., Johnson, F.E.: Social Responsibility of the Businessman. Harper, New York City (1953)
3. Bragdon, J., Marlin, J.: Is pollution profitable? Risk Manag. **19**(4), 9–18 (1972)
4. Cai, Y., Jo, H., Pan, C.: Doing well while doing bad? CSR in controversial industry sectors. J. Bus. Ethics **108**(4), 467–480 (2012)
5. Carroll, A.B.: Corporate social responsibility: evolution of a definitional construct. Bus. Soc. **38**(3), 268–295 (1999)
6. Cheng, B., Ioannou, I., Serafeim, G.: Corporate social responsibility and access to finance. Strateg. Manag. J. **35**(1), 1–23 (2014)
7. Dahlsrud, A.: How corporate social responsibility is defined: an analysis of 37 definitions. Corp. Soc. Responsib. Environ. Manag. **15**(1), 1–13 (2008)
8. Dyreng, S.D., Lindsey, B.P., Thornock, J.R.: Exploring the role delaware plays as a domestic tax haven. J. Financ. Econ. **108**(3), 751–772 (2013)
9. Eckbo, B.E., Nygaard, K., Thorburn, K.S.: Does gender-balancing the board reduce firm value? (2016)
10. Ferrara, E., De Meo, P., Fiumara, G., Baumgartner, R.: Web data extraction, applications and techniques: a survey. Knowl.-Based Syst. **70**, 301–323 (2014)
11. Johnston, K.A., Beatson, A.T.: Managerial conc1eptualisation of corporate social responsibility: an exploratory study (2005)
12. Laplante, S.K., Nesbitt, W.L.: The relation among trapped cash, permanently reinvested earnings, and foreign cash. J. Corp. Financ. **44**, 126–148 (2017)
13. Loughran, T., McDonald, B., Yun, H.: A wolf in sheeps clothing: the use of ethics-related terms in 10-k reports. J. Bus. Ethics **89**, 39–49 (2009)
14. Bromiley, P., Marcus, A.: The deterrent to dubious corporate behavior: profitability, probability and safety recalls. Strateg. Manag. J. **10**(3), 233–250 (1989)
15. Margolis, J.D., Elfenbein, H.A., Walsh, J.P.: Does it pay to be good? A meta-analysis and redirection of research on the relationship between corporate social and financial performance. Ann Arbor **1001**, 48109–1234 (2007)
16. Marr, C., Highsmith, B.: Tax holiday for overseas corporate profits would increase deficits, fail to boost the economy, and ultimately shift more investment and jobs overseas. Center on Budget and Policy Priorities, Washington, D.C. (2011)
17. Mills, R., et al.: WikiRate.org – leveraging collective awareness to understand companies' environmental, social and governance performance. In: Bagnoli, F., Satsiou, A., Stavrakakis, I., Nesi, P., Pacini, G., Welp, Y., Tiropanis, T., DiFranzo, D. (eds.) INSCI 2016. LNCS, vol. 9934, pp. 74–88. Springer, Cham (2016). doi:10.1007/978-3-319-45982-0_7
18. Moir, L.: What do we mean by corporate social responsibility? Corp. Gov.: Int. J. Bus. Soc. **1**(2), 16–22 (2001)
19. Namkung, Y., Jang, S.: Are consumers willing to pay more for green practices at restaurants? J. Hosp. Tour. Res. **41**(3), 329–356 (2017)
20. Parguel, B., Benoît-Moreau, F., Larceneux, F.: How sustainability ratings might deter greenwashing: a closer look at ethical corporate communication. J. Bus. Ethics **102**(1), 15 (2011)
21. SEC: Interactive data to improve financial reporting, June 2017. https://www.sec.gov/rules/final/2009/33-9002.pdf

22. Sethi, S.P.: A conceptual framework for environmental analysis of social issues and evaluation of business response patterns. Acad. Manag. Rev. **4**(1), 63–74 (1979)
23. Waddock, S.A., Graves, S.B.: The corporate social performance-financial performance link. Strateg. Manag. J. **18**(4), 303–319 (1997)
24. Wanderley, L.S.O., Lucian, R., Farache, F., de Sousa Filho, J.M.: CSR information disclosure on the web: a context-based approach analysing the influence of country of origin and industry sector. J. Bus. Ethics **82**(2), 369–378 (2008)
25. Wayne, L.: How delaware thrives as a corporate tax haven. N.Y. Times **30** (2012). http://www.nytimes.com/2012/07/01/business/how-delaware-thrives-as-a-corporate-tax-haven.html

Serendipity by Design? How to Turn from Diversity Exposure to Diversity Experience to Face Filter Bubbles in Social Media

Urbano Reviglio(✉)

LAST-JD International Joint Doctorate in Law, Science and Technology, University of Bologna, Bologna, Italy
urbanoreviglio@hotmail.com

Abstract. Personalization of online content creates filter bubbles and reinforces echo chambers. These are driven not only by natural human behaviours but also by design choices and efficiency-driven recommender systems. The traditional media policy goal of exposing citizens to diverse information to protect pluralism has not found its concrete application on social media. As the usage of social media as a news source increases, as well as personalization' sophistication and group polarization, there is a need for preventing audience fragmentation. The paper suggests serendipity as a potential design principle and, eventually, policy goal. Indeed, serendipity – considered both as a capability and a process of seeking and processing unexpected and valuable information – requires diverse information as a precondition and it causes cognitive diversity. Serendipity as a design principle might encompass fundamental phases of production and consumption of information, representing a positive freedom valuable from an epistemological, psychological and political perspective. With serendipity being both limited and cultivated in the digital environment, the research reveals a theoretical trade-off between relevance and serendipity (or unknown relevance) that might be tackled with serendipity-driven recommender systems and structural and informational nudging. Such approach could turn the media policy goal of exposing users to diverse information towards an experience of diversity that comes through an architecture for serendipity.

Keywords: Social media · Personalization · Filter bubbles · Serendipity · Media ethics

1 Introduction

Information online is increasingly essential for individuals to plan their lives rationally and to participate fruitfully in public life. In particular, the exposure to challenging or alternative information is fundamental for promoting critical thinking and informed decision-making, and even preventing or correcting inaccurate beliefs or dangerous radicalization. Though Internet gives unprecedented access to information and people of widely varied backgrounds, this increasing availability of diverse information does not guarantee an equally diverse exposure to different perspectives. Selective exposure

I. Kompatsiaris et al. (Eds.): INSCI 2017, LNCS 10673, pp. 281–300, 2017.
https://doi.org/10.1007/978-3-319-70284-1_22

to information – the individuals' tendency to favour consonant information and avoid dissonant one – can negatively affect deliberative discourse. Concerns increase when algorithms filtering provides highly personalised Internet experience, producing at the individual level *filter bubbles* and at the collective level *echo chambers*. Yet, whether these ultimately lead to collective political action or simply more fragmentation is a question that has been debated by social and political theorists ever since the Internet entered academia.

Media law and ethics discussed this issue in terms of "diversity exposure" without, however, being able to balance different ethical perspectives and offer an encompassing theoretical and methodological framework. Firstly, because there are not enough scientific evidences of the actual risks. Secondly, because exposing individuals to diverse information might polarize them even further, so it is not clear how much diverse information would be sufficient. At the same time, there are trends in social media's information production and consumption that are potentially limiting information diversity exposure. There is, indeed, a lack of autonomy in "de-personalising" the recommender systems that users enjoy. Social media are often opaque, especially about users' personal information consumption. Implicit personalization filtering is also increasing, as well as its accuracy. Furthermore, social media are designed to subtly affect users' behaviours in order to increase their engagement, and so the profit. Hence, individuals risk to be trapped in a continuous flow of highly personalised information rather than a more diverse and serendipitous one, without clear ways to opt-out, nor sufficient motivation, awareness and skills to do so. It is therefore becoming imperative to employ an "ethical mediator" to influence media users in an appropriate way.

The article explores this phenomenon through the lens of serendipity, considered both as a capability and a process of seeking and processing unexpected and valuable information. In fact, serendipity requires novel and diverse information as preconditions. Also, it causes cognitive diversity, creativity and innovation. Seen as a design principle, serendipity provides an alternative perspective that represents a positive freedom for users, as the best conceivable information encountering outcome, able to encompass fundamental phases of production and consumption of information, and valuable from an epistemological, psychological and political perspective.

Serendipity is, of course, a complex concept on which there is no consensus on the meaning. Apparently, it seems impossible both to create on demand, and to properly assess it. It will be necessary, then, to define what is meant with serendipity, how it expresses in the digital environment, and to explore whether it is possible to program it in algorithms, cultivate it by design, and objectively assess it. The major questions that the paper attempts to answer are the following: Is it possible to design serendipity? Can this be preferred to "diversity exposure" as a desirable media policy goal? Ultimately, can it ever be conceived as a useful normative horizon able to foster a healthier *infosphere* ecosystem [1]? If yes, how would that be possible?

In the first chapter, the paper aims at analysing through a literary review the main concerns about personalization, and whether one could legitimately argue for policy intervention. In the second chapter, the paper discusses the phenomenon of serendipity in the digital environment, and whether this might be considered an alternative design principle suitable to increase the diversity exposure in social media. It will be argued how this is possible by defining a preliminary framework of "serendipity by design". Finally, conclusions will be drawn.

2 The Risks of Hyper-personalised Information Filtering

By recommending content according to the inferred preferences of users, personalization performs a fundamental role of knowledge management. In general, it occurs explicitly and implicitly; the first exploits use of users' requests, while the other is mainly based on monitored users' activity. Both have increased dramatically in the last years, though many websites have acted to make passive forms of personalization the fastest growing forms [2]. Many fear that this hidden personalization could close users in a kind of cultural and ideological bubble [3, 4].

Over-personalization filtering can actually create filter bubbles and fuel echo chambers, two sides of the same token. The first is a kind of cultural and ideological bubble in which an individual continues to see, listen and read what reinforces its opinions and interests. The latter is a group situation where established information, ideas, and beliefs are uncritically spread and amplified, while dissenting views are ignored. The major difference between the two is that filter bubbles occur without the autonomy of the user [5]. Echo chambers, instead, pre-existed digital age: people naturally tend toward selective exposure and homophily [6].

Yet, there is little and conflicting research about the extent to which personalization is detrimental. This phenomenon, however, raises serious concerns both at the individual and collective level. Individual, because these filters might reduce opportunities for users' identity to self-determine. Indeed, personalization would reduce opportunities for serendipitous discovery, particularly reducing exposure to alternative points of view [7]. The consequences may be various: from the limitation of personal creativity, "insight and learning", to a reduction in our ability to build productive social capital [4]. Collective, because by fueling polarization, pluralism may be weakened and it may make people more vulnerable to censorship and propaganda or, to some extent, to self-propaganda. In fact, the filter bubble thesis is supported particularly in the context of online radicalisation [7, 8]. Furthermore, a frequent result is the formation of homogeneous, polarized clusters that are the primary driver of content diffusion, especially misinformation and conspiracy theories [9]. Finally, another prominent risk is growing inequality. Indeed, some privileged group of users, that have enough (digital) literacy, would be able to reach a good balance of information diversity, while a larger group of users are exposed only to a minimum, qualitatively inferior part of information [10].

The question is not whether filter bubbles and echo chambers exist because there is plenty of evidence of their existence [5, 8, 11–13]. As O'Hara and Stevens [14] argued, the two main questions are: whether social media's recommender systems are complicit in their growth and, in that case, whether this should be the target of a policy focus and, eventually, of intervention. Actually, it is problematic to answer these questions, since most of research is often inconsistent and rarely conclusive because it is generally survey-based, and thus dependent on self-reporting, or based on a small or unsatisfactory sample. Also, in the light of the rapidly changing media landscape, many studies become rapidly out-dated [15]. Currently, however, scholars agree on the idea that online news audiences are not more fragmented than offline news audiences [16].

Another fundamental issue is whether and how it is possible to dispel or weaken filter bubbles and echo chambers. In fact, some politically radical groups have claimed that they use exposure to opposition in order to strengthen their opinions further [8]. Simply exposing radicalised people to alternative points of view would not undo these phenomena [14]. Quattrociocchi and Sunstein [17] recently showed that debunking is ignored by $\sim$99.98% of conspiracy users on Facebook. In fact, for most of what people believe to know, they lack personal or direct information, so they must rely on what other people think. In some domains, this brings people to suffer from a "crippled epistemology" in the sense that they know very few things, and what they know is usually wrong [18]. Thus, extremism stems not from irrationality, but from the fact that they have little (relevant) information. It is a matter of *how* people acquire information. Hopefully, other studies [19–21] suggest that exposures to opposite views tend to increase political tolerance and awareness of the legitimate arguments underlying opposing perspectives. Indeed, individuals would strengthen their opinions after better understanding the other political side and, in some cases, formulating new opinions on issues that were against their former beliefs.

The possibility to counteract the potential risks of personalization is, therefore, controversial. Recommender systems and design choices, nonetheless, maintain a large role in preventing it. Hence, to some extent social media are complicit in the growth of filter bubbles. Though there is no empirical evidence that warrants any strong worries at present [15, 16], in the case that Artificial Intelligence would enable a sort of "hyper-personalization", that groups polarization would increase and that personalised news content would become people's main information source, problems for democracy could arise. These, actually, represent current trends, and prevention might be better than cure, even more in a fast-paced digital world.

2.1 Techno-regulation of Filter Bubbles by Design?

Internet constituted a significant challenge to traditional media policies and the role of public service media. With large media providers no longer serving a gatekeeping function, the diversity of individual exposure turned on the choices of individual users and algorithms. As a reaction to such new media environment, the notion of "diversity exposure" and alternative frameworks like "diversity by design" increasingly gained attention as a media policy [22]. These assume that pluralism is achieved when users actually enjoy a diverse media diet. Mere exposure to various sources and content, in fact, is not sufficient for ensuring actual users' experience of media diversity. A thorough analysis of media diversity in the digital age, therefore, should consider not only diversity supply and diversity exposure, but also the cognitive and affective factors that drive Internet use, and ultimately actual "diversity experience". Though most of the studies focused on policy implications at the macro level, the debate on pluralism must employ a user-centric perspective and thereby extend beyond the assumption that supply diversity equals diversity exposure [23].

The normative evaluation of such "diversity experience" is much more complex than for traditional diversity of supply [24]. Thus, fundamental questions for media ethicists and policy-makers arise: how to turn from diversity exposure to diversity experience? How can we conceptualise them as a possible policy goal? What are the values and goals that media diversity serves? And then: how much exposure to how many different contents and sources is considered sufficient? Is it possible to indirectly nudge users towards an enriching experience of diversity? These questions have hardly been discussed seriously in European media law and policy, even though filter bubbles are considered problematic for particular and different reasons for all models of democracy [13]. For Helberger [22] this can be partly explained by the uncertainty of policymakers about the permissibility of interference in an area that is as politically and legally sensitive as is individual information consumption. Certainly, users cannot be forced to choose diversely, and governments need to balance the fundamental freedoms of users and corporations.

The emergence of an "ambient intelligence" [25], in which automatic smart online and offline environments interact and take an unprecedented number of decisions in order to cater to our inferred preferences, represents a new paradigm in the construction of knowledge. Instead of merely pre-empting a user and providing her with inferred preferences, there should be a sustained interaction that engages a person as co-creator of a shared environment. Hence, lawyers and legislators should learn to articulate legal protection into the digitally enhanced environment, because human rights lack effective remedies if they are only articulated as written law [26]. Many scholars discussed what could be the role of institutions in designing such techno-regulation and, eventually, nudging people at least indirectly to expose and experience diverse information. Approaches like value-sensitive design (VSD) or pro-ethical design can help in this [27, 28]. It is possible, indeed, to nudge users to rely on 'empowerment nudges, which promote decision-making in the interests of citizens, as judged by themselves, without introducing further regulation or incentives' or manipulative measures [24, 29]. A similar approach is embraced by Floridi [28], which suggests what he calls "tolerant paternalism" that, in short, aims to modify the level of abstraction of the choice architecture by educating users to make their own critical choices and to assume explicit responsibilities. However, no concrete proposal has been made nor experimented yet.

The main rationale of any policy proposal that deals with the risks of hyper-personalization is that development of accurate beliefs requires the experience of diversity. On the one hand, selective exposure correlates with higher levels of attitudinal polarization and greater fragmentation in issue priorities whereas, overcoming the "majority illusion effect" [30], as the wrong perception of the popularity of certain belief or behaviours, may lead people to accept the legitimacy of disagreeable outcomes in the political sphere. Nonetheless, two main issues still raise doubts on the effectiveness of any preventive policy intervention; firstly, that likeminded media use can inspire greater rates of political participation. Yet, one may also argue that participation is not always good when it does not come from well-informed participants. And secondly, in some cases, counter-attitudinal information can even amplify attitudinal polarization, though in limited social groups. Complete solutions, therefore, must involve something more than mere exposure to alternative views.

3 Serendipity as a Design Principle

In the past, any procedure of information selection of information recognized and sought to solve in a beneficial manner the ideal tension between *relevance* – what a reader wants – and *serendipity* (or unknown relevance) – what a reader may want. In the digital environment, this balance inevitably shifted from serendipity to relevance [2]. Serendipitous information, indeed, represents any information that both users and algorithms still ignore that is actually relevant. Is it ever possible, then, to find out a balance in this theoretical trade-off, and to prospect a normative horizon able to cultivate a digital environment that can be freed – at least partially – from algorithms biases? Some scholars argue that in an age of user-driven pluralism, media policy goals and public service media could find new legitimacy in facilitating user experiences of diversity by creating encounters with surprising, challenging, and, ultimately, serendipitous content [23, 31], which can actually burst filter bubbles [32–36]. This is possible, for instance, by designing recommender systems that filter content that *intersect* user interests, or by nudging users with visualisation tools and design choices in order to better manage their information diet – what I will refer to as "meta-personalization". Preliminary proposal to the design of such prospective principle are presented in this chapter by discussing the most promising software, research and visualization tools. Before that, it is fundamental to look deeper into the nature and the dynamics of serendipity. By doing this, it will first be analyzed the main characteristics of such phenomena. Then, it will be introduced a brief debate about serendipity in the digital environment and, finally, a preliminary framework to "serendipity by design", its expected outcomes and unintended consequences.

3.1 Origins and Outcomes of Serendipity

The fascinating concept of serendipity originated from the Persian fairy tale "The Three Princes of Serendip", which narrates how these traveling princes were "always making discoveries, by accidents and sagacity, of things which they were not in quest of" [37]. Serendipity, indeed, is the art of discovering new things by observing, and learning from unexpected situations. It can be broadly defined as "an unexpected experience prompted by an individual's valuable interaction with ideas, information, objects, or phenomena" [38]. Another useful definition considers serendipity as "the finding of something valuable without its being specifically sought" [39]. This definition, indeed, encapsulates serendipity in the act of finding, while at the same time denoting the necessity of value, distinguishing itself from coincidence, luck and randomness. In fact, the discovery can be perceived as accidental but not necessarily unplanned or the result of fortuity. More often serendipity is the result of a lot of groundwork, observation, and knowledge. As Louis Pasteur once famously said [37]: "in the field of observation, chance favours only the prepared mind". Thus, serendipity is also a capability [40]. Aside from mainstream interpretations, serendipity has a dual nature [41], as it is "neither only a process nor only a perception, but rather has a paradoxical nature dependent on conditions and context". It is true, indeed, that serendipity is not a buzzword *tout court* but it has different dimensions, environmental factors and components [42], researched by "serendipity studies". Instead, for a long time media researchers largely ignored the role of serendipity in news discovery [43], while recently got momentum.

Serendipity plays a fundamental role in our everyday life. Yet, the ability to extract knowledge from an unexpected event covers all areas of human activity, including business, law, politics and, particularly, science. According to sociologist Robert K. Merton [36] serendipity is the "happy accident" inherent in scientific research, one of the main forces that has steered the progress of science. It has been estimated that over 50% of scientific discoveries were unintended and it is full of notorious examples [44]. Hence, its role on epistemology of science is well established, even though there is a lack of advancement in this field. Campanario [45] suggests the reason for this in the mythisation of scientific research. Science is supposed to be something that is totally under control, so that scientists may be reluctant to admit that the discoveries for which they are honoured were accidental. To some extent this conjecture can be advanced to our own illusion to have full control over our lives. Though relatively less momentous, serendipitous insights follow also in our everyday lives, and can change our day, or even our life. Indeed, serendipity helps us to innovate and to be creative, leading us to the emergence of a theory, a law or perhaps simply to revise an opinion, which had never been planned and therefore not intentionally sought for [45]. While not every creative moment is necessarily serendipitous, all serendipitous moments are creative ones [46]. In fact, serendipity has been considered as a fundamental experience to maintain creativity in the computer era [47].

In the discussion about Internet developments there have been concerns about the loss of serendipity as well as the acknowledgment of its fundamental value [34]. In *Being Digital* [48] Nicholas Negroponte already anticipated the possibility to crank personalization up or down. He argued that "depending on time available, time of day, and our mood, we will want lesser or greater degrees of personalization" to have "a slider that moves both literally and politically from left to right to modify stories about public affairs" because "we may wish to experience the news with much more serendipity". Then, in 1997 Ted Gup [49] heralded *"The End of Serendipity"*, warning that the tools of efficiency-quick retrieval were helping us to better find information we need, but hindering us from accidentally discovering information we need, but did not realize we needed. Eventually, recommender systems improved, as did the recognition of the need for serendipity and its preconditions.

Sociologist Abbott [50] provided a theoretical basis for such concerns arguing the necessity to defend randomness from processes for search and discovery. Meckel [51] even made an appeal to literally "save serendipity", a fundamental experience for our identities undermined by profiling algorithms that "are forever stuck in the past, as they base their calculations on our actions in times foregone" so that they may "force us into a never-ending time-warp, dwelling forever in the status quo of our own preferences and desires." The value of serendipity has been stressed also in more specific contexts. MIT Media Lab director Ethan Zuckerman is a well-known exponent of serendipity, in particular concerning the development of a cosmopolitan culture and how its development is limited due to georeferenced algorithms [52]. Also, the legal scholar Cass Sunstein [3], a leading expert in political polarization, in his third version of *Republic 2.0* argues that his "largest plea […] is for an architecture of serendipity" because "chance encounters and shared experiences" are preconditions for a well-functioning democracy, and serendipity would sustain them and counteract homophily, and promote both self-government and individual liberty. Then, he proposes that Facebook could create a "serendipitous news and opinions button, allowing people to opt in". Thus, it is already a widespread opinion that serendipity is a valuable, if not fundamental, experience in the digital environment, especially given its technical volatility.

Actually, relatively little research has been undertaken to assess how well existing and novel approaches to information interaction support serendipity [38]. Most of the studies on serendipity are indeed conducted perforce retrospectively [53]. However, new promising experimentations are digging the ground [34]. The study of serendipity in information studies and computer science is an emerging research area, in particular for Human Computer Interaction (HCI) [54]. Of course, one may argue that Internet is already very serendipitous. Computer itself has been considered a "serendipity machine" [55]. Hyperlinked digital environments are a fertile ground for serendipity [36], providing a diversity of resources to which users may not have otherwise been exposed [56]. Also, social media are a powerful source of coming across information serendipitously [57]. Indeed, it is often a tacit design goal for engaging users. Hence, at first sight a loss of serendipity in the digital environment could seem only a philosophical speculation. Serendipity in the digital environment, however, is very different from the "offline serendipity", as it can be effectively both cultivated and limited not only by recommender systems, but also with affordances [58], including software, visualization tools and design choices (as it will be explored in the next chapter). It is, therefore, a matter of degree, and how information societies will be able to reach a desirable informational balance. Social media, in particular, can indeed be ever more serendipitous.

Also, the value of serendipity must be considered in light of current political landscape and potential technological developments. On the one hand, as populism is currently growing across both Europe and the United States in a period of perpetual crisis, there are concerns that more and more people who were previously confident enough to voice their beliefs openly, thanks to social media are connecting to like-minded others and becoming more confident and vigorous. On the other hand, future advances in artificial intelligence and the semantic web have the potential to enable algorithms to make ever more sophisticated recommendations, while virtual reality, augmented reality and the Internet of things will definitely blur the distinction between online and offline. As Domingos [59] argues, tomorrow's cyberspace will be a vast parallel world that selects only the most promising information to try out in the real one, and it will be like "a new, global subconscious, the collective id of the human race". Designing for serendipity might become valuable in other emerging contexts. In the near future, more and more accurate decisions will be delegated to algorithms, design for serendipity might ensure a way to opt-out and, eventually, to opt-in to seek the unexpected.

Media and information ethicists and policy makers need to ask not only whether and how Internet users need to be sustained to seek new, diverse, challenging and, ultimately, serendipitous information, but, more importantly, also how to translate abstract theoretical concepts in a concrete methodological framework. In other words, how to avoid that such potential principle would remain only a utopian one?

3.2 "Serendipity by Design"

Programming for serendipity sounds like an oxymoron. In fact, how can a subjective process that is only discernible in hindsight and for which the unexpected plays such an important role in its perception be supported? Indeed, serendipity cannot be created on demand. Nonetheless, it can be cultivated by creating opportunities for it through the design of physical, digital, and learning environments [34]. Broadly speaking, from a user-perspective there are two information behaviours for seeking serendipity: non-purposive or passive (like browsing recommender systems selection) and purposive or active (like nudging with visualisation tools and design choices) [35]. Similar to other proposals to expose users to different media, like "diversity by design" [22], serendipity by design is the idea that it is possible to create an architecture that helps people to make diverse choices and to give the incentives for seeking alternative information and, ultimately, to have more serendipitous insights, which can represent the best conceivable outcome from a user perspective.

In the following chapter, it will be briefly presented how serendipity might be designed and cultivated. First, it will be assessed the algorithmic metrics underpinning digital serendipity in recommender systems (even though it is not the specific expertise of the author). Secondly, it is explored how users might personalize their personalized experience. Finally, limitations and unintended consequences will be briefly discussed.

3.2.1 Serendipitous Recommender Systems

The primary goal of recommender systems (RSs) is to provide personalized recommendations so as to improve users' satisfaction. Many studies are increasingly showing that RSs are moving beyond accuracy and embracing serendipity [60]. Indeed, user satisfaction does not always correlate with high recommender *accuracy/precision*. In fact, when we successfully implement serendipitous encounters in RSs, it is possible to avoid predictable recommendations in collaborative filtering systems, to solve the over-specification problems in content-based systems and also to help users reveal their unexpected interests [61]. Nevertheless, most social media's RSs aim to instant gratification more than long-term satisfaction. Indeed, most of the current recommender systems have been criticized to not sufficiently account for serendipity [33]. Studies have also showed that users are willing to sacrifice some amount of accuracy for improved serendipity in the algorithms filter performance [62]. In some cases, RSs metrics such as novelty, diversity and serendipity can also be improved simultaneously, without any apparent trade-off between the three qualities [63]. A question, therefore, arises: does accuracy as a RS metric naturally lead to more profits for social media rather than serendipity-driven RSs? In other words, it seems that very accurate personalization increases users' engagement in the short-run more than might serendipity which inevitably involves more chance in its design. Arguably, this is an emerging trade-off in knowledge management (Table 1).

Table 1. A theoretical trade-off in recommender systems.

Platform	Main values	Explicit goal	Latent goal	Tendency
Business Social Media	Relevance	Engagement (short-term gratification)	Profit	Determinism
"Alternative Social Media"	Serendipity	Unexpected discoveries (long-term satisfaction)	Individual and social fulfillment	Chance and randomness

Actually, there is no consensus on the definition of serendipity in RSs [34]. However, the two core characteristics of serendipity embedded in RSs are usually *unexpectedness/surprisingness* and *usefulness/relevancy* [61, 63]. Serendipity builds also upon the concept of *novelty*, but expands this by the factor of a positive discovery, because unexpected and useful [64]. Further, other metrics are underpinned by serendipity such as *diversity* and *unfamiliarity*. This is certainly relevant in the context of media policy's attempts to cultivate information diversity exposure and experience. Serendipity can encompass diverse metrics: a user is surprised by a novel, unexpected (thus diverse) and relevant information, see Fig. 1.

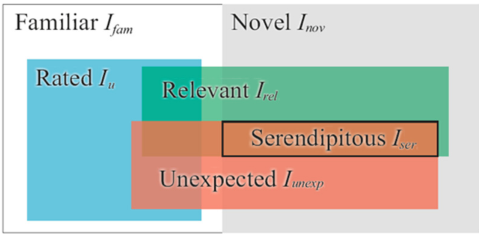

Fig. 1. Recommended items from a user's point of view (Kotkov et al., 2016).

Of course, optimizing current RSs for more serendipitous recommendations is not a trivial task. Firstly, increasing serendipity might negatively affect accuracy. Therefore, it has to be strategically done in order to alleviate the risk of having a distrust effect [32]. Hence, a good balance between novelty, serendipity and accuracy of recommendations seems to be necessary. There have already been, nonetheless, promising attempts of stimulating serendipitous insights in RSs, like the "serendipity equation" by Campos and Figueireido [65–67]. Notably, Max, their information retrieval software, reached the level of 52.7% of (pseudo)serendipitous suggestions. Despite the subjective character of the results, it seems that programming *for* serendipity is indeed possible. Among dozens, or even hundreds, of variables, it is then possible to design serendipity-driven algorithms. Interestingly, in their evaluation of Max, Campos and Figueiredo offered a granular approach to the question of assessing serendipitous findings and encountering: by dividing user results into six distinct categories, according to their possible outcomes. By opening space for similar granularity in other approaches to digital interactions, it is possible to expand and encompass a wide range of results that can be considered valuable. In fact, under the design commitment to serendipity there are several other concepts related to information seeking, for instance: pseudo-serendipity; opportunistic acquisition of information [68]; micro-serendipity [69]; accidental discovery of information [70] etc. Designing for serendipity would indeed mean to aim to cultivate all of these indirectly, and openly serendipity par excellence.

Other search environments have been designed with the purpose of creating opportunities for accidental information discovery [13, 34]. Paradigmatically, some of these are usually effective in supporting serendipity, but the recommendations are not always considered particularly interesting or relevant. This unveils again a trade-off related to the previous one: the more you want to discover accidental information, the more you need to invest your time in seeking it (because filtering is less accurate). Of course, looking for serendipity is time-consuming. This highlights the importance of considering the context in which users may be searching; accidental information discovery may be more useful when trying to get news of developments in a specific area rather than when trying to get a background understanding. Otherwise, in the perspective of Negroponte [48], sometimes one may want to access the personalized *"daily me"*, and some other times the more generalized *"daily us"*. In essence, the "sweet spot" is likely to differ based on the user's information needs and tasks.

Actually, users have not much influence on the recommendation process itself, apart from providing implicit or explicit ratings for items, usually in the long-term use of the system [71]. Several problems indeed arise from the limited degree of interactivity and user control over the recommendation process. For instance, a lack of transparency prevents users from comprehending why certain items are suggested and, as a potential cause, a reduced trust in the system is expected. In fact, users often desire a more active role in the recommendation process. Interactive control might increase the system's transparency and acceptance, and the design for serendipity actually implies it.

From a policy perspective, it seems that compel social media to optimize RSs for serendipity does not appear as a viable approach. As one can imagine, its implementation would be extremely problematic, not to say that consensus on the metrics has to be reached. Instead, from an ethical perspective, such approach can be easily considered a form of stealthy paternalism. Furthermore, the aim of serendipity by design would be to find a right balance between personalization (relevancy-driven recommendations) and personalized 'generalization' (serendipity-driven recommendations). It is their dynamic relation which can better sustain a desirable serendipitous environment. Therefore, serendipity-driven RSs can better represent a principled commitment. What is probably more viable, then, is a structural and informational nudging that stimulate the activity of users.

3.2.2 Serendipitous Meta-personalization

Aside from RSs, it is possible to design several potential features that could nudge users to seek and encounter diverse, challenging and, eventually, serendipitous information [38]. Opening the possibility for users to personalize their personalized filtering experience is what one may call "meta-personalization", namely dynamically personalize personalization outcomes depending on needs, tasks and mood. Considering the widespread lack of users' interactive control and awareness about their information consumption, there are indeed several ways to empower users without being particularly intrusive nor radically against social media's policies; extracting value from their profiles meaningfully, illustrating connections, and to stimulating creative and serendipitous associations.

Speculatively, fundamental features for a more serendipitous environment in social media might be: firstly, visualization tools about users' *information diet* (awareness); secondly, *multiple filtering*, namely to browse diverse filters per person (interactive control) and, thirdly, an architecture that nudges users towards a more diverse information exposure, like Sunstein's proposal of a serendipitous button to encounter challenging political content (soft nudging).

First of all, to stumble upon more serendipitous information awareness of users' information diet is essential. To mention a few software that might be implemented by default, the visualisation design developed by Nagulendra and Vassileva [72, 73] displays to users their filter bubbles. By showing them which categories and friends are in their bubble and which ones are not. Users can simply control the algorithm by manipulating the visualisation to escape the bubble, by adding or removing friends on a certain topic to the filters. The results are truly promising: 72% of participants said that it was easy to find an interest which was not inside their filter bubble, so that they were able to "discover new interests that they didn't display otherwise in their behaviour" [73]. Then, *Balancer* shows an approximate histogram of the user's liberal and conservative pages, with the aim to increase user-awareness so that they would make their reading behaviour more balanced. A similar aim is taken by *Scoopinion* which visualize what journals you read and for how long. Such tools also aim to increase the transparency of information intermediaries and shed the light on user information consumption. They can nudge users to shift towards more balanced reading behaviour, though it can also lead some users to more polarized or unbalanced views [74]. Of course, given that users' proclivities vary among individuals and topics, effects of certain interface designs on users' behaviour may also vary. Nevertheless, a multi-layered serendipitous architecture could favour more diversity experience in the majority of users.

Furthermore, a fundamental affordance to cultivate a serendipitous environment might occur by overcoming the dogma of one profiling per user. In fact, in mainstream social media there is no possibility to have more identities (namely more filtering). As Facebook Chief Mark Zuckerberg itself believes: "having two identities for yourself is an example of a lack of integrity" [75]. Instead, identity is a dynamic experience, so that one may want (and need) to better adapt RSs, and seek different information at a different time. In fact, many factors, such as weather, mood or location, can influence user preferences for recommended items [76]. Changing the current approach can actually increase users' resiliency in seeking more serendipitous information by subtracting from the determined path offered by personalization [5].

Ideally, any user might be able to access information filtered in a varied spectrum from the most relevant to the most potentially-serendipitous, namely from relevant personalized (daily me), through generalized (daily us), to random-serendipitous (by chronological order which, for example, is already available in Facebook but it resets itself after a couple of hours). As shown in Fig. 1, these would actually satisfy the three broad information environments that sustain the information-seeking processes that lead to serendipity [77]. The first, design-centred information environment, is based on interests and users' preferences (like the "daily me"). Then, user-centred information environment is intended to address a broad community of users' needs rather than the designers' needs (like the "daily us"). And finally, an environment in which order is absent (i.e. by chronological order) (Fig. 2).

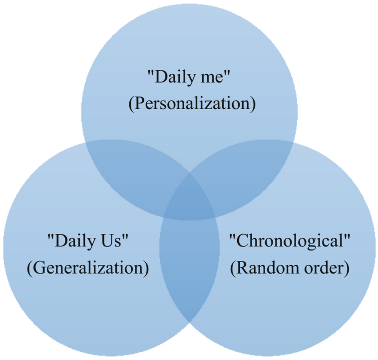

Fig. 2. An information environment that sustain the information-seeking processes that lead to serendipity.

Yet, many other design choices can be explored to increase the exposure to diverse information. For instance, as multiple personalized filters could be deployed, any user could even re-start its profiling from scratch. Also, one may argue for an "informative empathization", namely the possibility to browse (at request) the algorithmic filtering of others' users/friends. Even more importantly, graphs could show users whether their exposure is "sustainable", and provide information diet *suggestions* for a good balance between daily me and daily us, based on shared parameters (I.e. EU-based). These can make users aware of the need for diverse information. At the same time, it would be as important to provide the possibility to *filter the kind of the content*, such as articles, friends' lives, videos, opposing political.

In conclusion, if one analyzes the emerging phenomenon of hyper-personalization through the lens of serendipity, it is possible to prevent that and rebalance many of the well-established rights of information intermediaries. Thus, any user would eventually benefit from having the autonomy to meta-personalize their information seeking process, effectively access all the information of its own social network, and enjoy an information diet analysis through interactive visualization tools. Taken all together, such design choices might represent a structural and informational nudging approach which is not necessarily intrusive. Certainly, it is easy to consider such requests as mere speculations with no concrete potential for being applied, because it is based on assumptions without specific boundaries that could actually permit almost any sort of requests. However, as long as the research would increase and demonstrate the beneficial outcomes of designing serendipity, it could be possible to establish a methodological framework of "serendipity by design", and to find compromises for any eventual implementation. Indeed, this might become the proactive commitment and, eventually, the actual legitimization of public-service media in the digital age.

3.2.3 Limitations and Unintended Consequences

There are several limitations and potentially unintended consequences for such preliminary proposals that have been considered so far. These can be roughly summed up as ineffectiveness – most people indeed won't use such affordances and some other would even personalize further their social media experience – and unfeasibility – social media would never come to such a compromise, and any assessment of serendipity will never be objective and scientific.

Firstly, is it really sufficient to provide the autonomy for a "serendipitous personalized diversity exposure" to fulfil the pluralistic media goal to expose citizen to information diversity? Of course, it is not. Studies on de-polarizing techniques are more often contrasting. Usually some people polarize further, while others enjoy more information diversity. Even worse, it seems that people that ideally would benefit more from being exposed to challenging perspectives are those less willing. Hence, for some people it is plausible that design serendipity becomes only a means to further personalize their Internet experience. This, however, does not represent a valid argument against such potential principled framework. Research must establish benefit for the majority of the users. Then, it is obvious that the answer lies in education in the first place. In this respect, it is worth acknowledging that there is also room to study and teach serendipity as a capability in educational settings [47]. And as educators might not be available during networked learning, a mixed supply of information provided through the mediation of people could advance self-directed networked learning [78]. Usually, the challenges that users must overcome to experience media diversity in an online environment and, eventually serendipity, are rather well known [22]. These must be taken into consideration when designing for serendipity.

Then, another prominent issue is how to frame a measure for the assessment of perceived serendipity afterwards. It is perhaps one the most problematic issue. Foremost, a clear theoretical and methodological framework of 'serendipity by design' it is needed. In particular, the boundaries of the human rights underpinned by serendipity need to be established, and then how this could be assessed can be approached. There are, nonetheless, already experimented techniques. Apart from indirect assessment throughout surveys and other new methodologies [69], it is possible to observe serendipity in controlled research environments, by directly observing information encountering behaviour rather than relying only on self-reported data [35], for instance with the "critical incident techniques" on representative samples [79], or more specifically, Information Seeking in Facebook Scale (ISFS) [80]. However, in order to quantify serendipity at a large scale, a principled method to automatically identify serendipitous events is needed and has still to be defined [81]. Additional measures for assessing the effectiveness of "serendipity by design" might be taken into consideration, like the percentage of people changing idea about political issues in social media.

To conclude, from a media ethics perspective this preliminary framework might convey the principle of media pluralism, updated to the emerging practices of media exposure. Serendipity as a media goal would convey different perspectives in a more general and, essentially, positive goal. Arguably, even more serendipitous and personalized information diversity exposure in a more dynamic and interactive environment, could be considered sufficient outcomes in the age of personalization. Even if such architecture of serendipity might be even much more complex to achieve than "diversity experience", it represents a better goal, both from an epistemological and a user-centric perspective. Such design architecture would not represent a complete solution but, nonetheless, it has the potential to prevent information redundancy, and further audience fragmentation. Of course, any concrete form of policy intervention with the goal of promoting it will require a more systematic and detailed conceptualisation of serendipity and its design, because little experience exists so far.

4 Conclusions

The paper analyzed serendipity as an alternative design principle for social media. With "serendipity by design" was intended such alternative approach to systems engineering which takes serendipity into account throughout the design process and the interaction of users within the system. The ultimate goal is to cultivate an even more serendipitous environment in order to trigger more serendipitous insights, as the best conceivable outcome from a user perspective. Serendipity, in fact, is a fundamental feature of social reality, a complex process which occurs in all realms of human life, especially in scientific progress. It represents an inherently positive experience that embraces fundamental values and experiences like curiosity, research, discovery, and identity development. It can, however, be limited and facilitated by design choices and recommender systems. In particular, a potential hyper-personalization of online content could become increasingly deterministic, and serendipity might introduce more randomness to prevent dystopic scenarios. Thus, it has been argued that the experience of serendipity is not in danger but it has to be preserved and cultivated.

Seen in its paradoxical nature of being both a capability and an accidental perception, thus dependent on conditions and context, serendipity is suitable to encompass fundamental phases of production and consumption of information: from algorithmic filtering to information behaviour capabilities. Considering it as a design principle helps to stress how it can tackle the potential risks of hyper-personalization and redesign the media policy protection of pluralism by strengthening individual self-determination. To occur, in fact, serendipity requires novel and diverse information as preconditions. Eventually, it also causes cognitive diversity. Such aims might be facilitated by designing recommender system that are serendipity-driven that, broadly speaking, intersect users' interests and, also, by providing further innovative and interactive affordances to users.

The presented framework is also meant to imply more autonomy and awareness to increase the capability of users to seek, encounter and process more serendipitous information – what has been called meta-personalization. In particular, the potential consequences to unfold multiple filtering and to provide an effective information diet management has been stressed. Of course, these represented arbitrary affordances among many possible, though essential ones to cultivate a more dynamic environment where diverse information can be more easily accessed. These could indeed shuffle the cards of information management power, bringing it back to the users. Yet, it is questionable to what extent would users be willing to actually experiment with them. Digital literacy and critical thinking remain fundamental *condition sine qua non* for a beneficial expression of serendipity. Serendipity as a capability, however, can represent an educational goal for individuals to flourish in mature information societies. It can, in fact, increase open-mindedness, creativity and innovation. As it arises more often through inter-disciplinary scholars, it helps to bridge ideas out of users' information bubble consumption. Overall, serendipity could actually turn the media policy goal of expose users to diverse information towards an experience of diversity that comes through its design, both in digital environments and in individuals' information behaviour.

Far from having systematically framed its boundaries, a preliminary framework for what has been termed "serendipity by design" has been presented. Such approach has the potential to help designing systems and nudging individuals to be more resilient to the redundancy coming from increasingly sophisticated content personalization. To some extent, filter bubbles would probably persist but, eventually, users could burst them and autonomously and effectively seek and experience diverse information through an architecture that sustains more potentially-serendipitous encounters. Certainly, given the extraordinary complexity of social media, its rapid evolution, and potential unintended consequences, some realism about the above-mentioned outcomes must be maintained. Yet, there is room for further experimenting the potential of designing serendipity in social media.

Acknowledgements. This research is funded by the ERASMUS MUNDUS program LAST-JD, Law, Science and Technology coordinated by University of Bologna.

References

1. Floridi, L.: The Fourth Revolution: How the Infosphere is Reshaping Human Reality. OUP Oxford, Oxford (2014)
2. Thurman, N.: Making 'The Daily Me': technology, economics and habit in the mainstream assimilation of personalized news. Journalism **12**(4), 395–415 (2011)
3. Sunstein, C.R.: # Republic: Divided Democracy in the Age of Social Media. Princeton University Press, Princeton (2017)
4. Pariser, E.: The Filter Bubble: How the New Personalized Web is Changing What We Read and How We Think. Penguin, Westminster (2011)
5. Bozdag, E., Timmermans E.: Values in the filter bubble ethics of personalization algorithms in cloud computing. In: 1st International Workshop on Values in Design–Building Bridges between RE, HCI and Ethics (2011)

6. Lazarsfeld, P.F., Berelson, B., Gondel, H.: The People's Choice, New York, Duell, Sloan, Pearce (1944)
7. Sunstein, C.R.: Going to Extremes: How Like Minds Unite and Divide. Oxford University Press, Oxford (2009)
8. Wojcieszak, M.: 'Don't talk to me': effects of ideologically homogeneous online groups and politically dissimilar offline ties on extremism. New Media Soc. **12**(4), 637–655 (2010)
9. Vicario, M., Bessi, A., Zollo, F., Petroni, F., Scala, A., Caldarelli, G., Quattrociocchi, W.: The spreading of misinformation online. Proc. Natl. Acad. Sci. **113**(3), 554–559 (2016)
10. Gazoia, A.: Senza filtro. Chi controlla l'informazione. Minimum Fax (2016)
11. Wojcieszak, M.E., Mutz, D.C.: Online groups and political discourse: do online discussion spaces facilitate exposure to political disagreement? J. Commun. **59**(1), 40–56 (2009)
12. Bakshy, E., Messing, S., Adamic, L.A.: Exposure to ideologically diverse news and opinion on Facebook. Science **348**(6239), 1130–1132 (2015)
13. Bozdag, E., van den Hoven, J.: Breaking the filter bubble: democracy and design. Ethics Inf. Technol. **17**(4), 249–265 (2015)
14. O'Hara, K., Stevens, D.: Echo chambers and online radicalism: assessing the internet's complicity in violent extremism. Policy Internet **7**(4), 401–422 (2015)
15. Zuiderveen Borgesius, F.J., Trilling, D., Moeller, J., Bodó, B., De Vreese, C.H., Helberger, N.: Should we worry about filter bubbles? Internet Policy Rev. J. Internet Regul. **5**(1), 1–16 (2016)
16. Fletcher, R., Nielsen, R.K.: Are news audiences increasingly fragmented? a cross-national comparative analysis of cross-platform news audience fragmentation and duplication. J. Commun. (2017)
17. Quattrociocchi, W., Scala, A., Sunstein, C.R.: Echo chambers on Facebook. In: Harvard Law Review (2016)
18. Sunstein, C.R., Vermeule, A.: Conspiracy theories: causes and cures. J. Polit. Philos. **17**(2), 202–227 (2009)
19. Xing, X., Meng, W., Doozan, D., Feamster, N., Lee, W., Snoeren, Alex C.: Exposing inconsistent web search results with bobble. In: Faloutsos, M., Kuzmanovic, A. (eds.) PAM 2014. LNCS, vol. 8362, pp. 131–140. Springer, Cham (2014). doi:10.1007/978-3-319-04918-2_13
20. Semaan, B.C., Robertson, S.P., Douglas, S., Maruyama, M.: Social media supporting political deliberation across multiple public spheres: towards depolarization. In: Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing, pp. 1409–1421. ACM (2014)
21. Liao, Q.V., Fu, W.T.: Can you hear me now?: mitigating the echo chamber effect by source position indicators. In: Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing, pp. 184–196. ACM (2014)
22. Helberger, N.: Diversity by design. J. Inf. Policy **1**, 441–469 (2011)
23. Hoffmann, C.P., Lutz, C., Meckel, M., Ranzini, G.: Diversity by choice: applying a social cognitive perspective to the role of public service media in the digital age. Int. J. Commun. **9** (1), 1360–1381 (2015)
24. Helberger, N., Karppinen, K., D'Acunto, L.: Exposure diversity as a design principle for recommender systems. In: Information, Communication & Society, pp. 1–17 (2016)
25. Hildebrandt, M., Koops, B.J.: The challenges of ambient law and legal protection in the profiling era. Modern Law Rev. **73**(3), 428–460 (2010)
26. Hildebrandt, M.: Legal protection by design: objections and refutations. Legisprudence **5**(2), 223–248 (2011)
27. Friedman, B.: Value-Sensitive Des. Interact. **3**(6), 16–23 (1996)
28. Floridi, L.: Tolerant paternalism: pro-ethical design as a resolution of the dilemma of toleration. Sci. Eng. Ethics **22**(6), 1669–1688 (2016)

29. Helberger, N., Klein-von Königslöw, K., van der Noll, R.: Convergence, information intermediaries and media pluralism—Mapping the legal, social, and economic issues at hand. A quick scan. Institute for Information Law, Research report (2014)
30. Lerman, K., Yan, X., Wu, X.-Z.: The “majority illusion” in social networks. PLoS ONE **11**(2), e0147617 (2016). doi:10.1371/journal.pone.0147617
31. Helberger, N.: Public service media merely facilitating or actively stimulating diverse media choices? public service media at the crossroad. Int. J. Commun. **9**, 17 (2015)
32. Semaan, B.C., Robertson, S.P., Douglas, S., Maruyama, M.: Social media supporting political deliberation across multiple public spheres: towards depolarization. In: Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing, pp. 1409–1421. ACM (2014)
33. Matt, C., Benlian, A., Hess, T., Weiß, C.: Escaping from the Filter Bubble? The Effects of Novelty and Serendipity on Users’ Evaluations of Online Recommendations (2014)
34. Race, T., Makri, S.: Accidental Information Discovery: Cultivating Serendipity in the Digital Age. Chandos Publishing, Oxford (2016)
35. Kotkov, D., Wang, S., Veijalainen, J.: A survey of serendipity in recommender systems. Knowl.-Based Syst. **111**, 180–192 (2016)
36. Corneli, J., Pease, A., Colton, S., Jordanous, A., Guckelsberger, C.: Modelling serendipity in a computational context (2014)
37. Merton, R.K., Barber, E.: The Travels and Adventures of Serendipity: A Study in Sociological Semantics and the Sociology of Science. Princeton University Press, Princeton (2006)
38. McCay-Peet, L., Toms, E.G.: Proposed Facets of a Serendipitous Digital Environment (2013)
39. Boden, M.A.: The Creative Mind: Myths and Mechanisms. Psychology Press, Hove (2004)
40. de Rond, M.: The structure of serendipity. Cult. Organ. **20**(5), 342–358 (2014)
41. McBirnie, A.: Seeking serendipity: the paradox of control. In: Roberts, S.A. (Ed.), Aslib Proceedings, vol. 60(6), 600–618. Emerald Group Publishing Limited (2008)
42. Edward Foster, A., Ellis, D.: Serendipity and its study. J. Doc. **70**(6), 1015–1038 (2014)
43. Wright, C.: Stumbling and sharing: smartphones and serendipity in online news encounters. University of Missouri-Columbia (2015)
44. Dunbar, K., Fugelsang, J.: Scientific thinking and reasoning. In: The Cambridge Handbook of Thinking and Reasoning, pp. 705–725 (2005)
45. Campanario, J.M.: Using citation classics to study the incidence of serendipity in scientific discovery. Scieontometrics **37**(1), 3–24 (1996)
46. Melo, R., Carvalhais, M.: Regarding value in digital serendipitous interactions. J. Sci. Technol. Arts **8**(2), 37–44 (2016)
47. Sawaizumi, S., Katai, O., Kawakami, H., Shiose, T.: Using the concept of serendipity in education (2007)
48. Negroponte, N.: Being digital. Vintage, New York (1996)
49. Gup, T.: Technology and the end of serendipity. Chronicle High. Educ. **44**(21), A52 (1997)
50. Abbott, A.: The traditional future: a computational theory of library research. College Res. Libr. **69**(6), 524–545 (2008)
51. Meckel, M.: “SOS – SAVE OUR SERENDIPITY”, Personal Blog, 11 October 2011. http://www.miriammeckel.de/2011/10/11/sos-save-our-serendipity/?goback=%2Egde_1975471_member_112072885
52. Zuckerman, E.: Rewire: digital cosmopolitans in the age of connection (2013)
53. Gabriel, Y., Muhr, S.L., Linstead, S.: Luck of the draw? serendipity, accident, chance and misfortune in organization and design (2014)

54. Makri, S., Blandford, A.: Coming across information serendipitously - Part 1, Process Model J. Doc. **68**(5), 684–705 (2012)
55. Green, D.: The Serendipity Machine. A Voyage of Discovery Through the Unexpected World of Computers. CreateSpace, Paramount (2008)
56. Thurman, N., Schifferes, S.: The future of personalization at news websites: lessons from a longitudinal study. J. Stud. **13**(5–6), 775–790 (2012)
57. Dantonio, L., Makri, S., Blandford, A.: Coming across academic social media content serendipitously. Proc. Am. Soc. Inf. Sci. Technol. **49**(1), 1–10 (2012)
58. Gibson, J.J.: The Ecological Approach to Visual Perception: Classic Edition. Psychology Press (2014)
59. Domingos, P.: The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World. Basic Books, New York (2015)
60. Ziegler, C.N., McNee, S.M., Konstan, J.A., Lausen, G.: Improving recommendation lists through topic diversification. In: Proceedings of the 14th International Conference on World Wide Web, pp. 22–32. ACM (2005)
61. Ge, M., Delgado-Battenfeld, C., Jannach, D.: Beyond accuracy: evaluating recommender systems by coverage and serendipity. In: Proceedings of the Fourth ACM Conference on Recommender Systems, pp. 257–260. ACM (2010)
62. Zhang, Y.C., Séaghdha, D.Ó., Quercia, D., Jambor, T.: Auralist: introducing serendipity into music recommendation. In: Proceedings of the Fifth ACM International Conference on Web Search and Data Mining, pp. 13–22. ACM, ISO 690 (2012)
63. André, P., Teevan, J., Dumais, S.T.: Discovery is never by chance: designing for (Un) serendipity. In: Proceedings of the Seventh ACM Conference on Creativity and Cognition, pp. 305–314. ACM (2009)
64. Herlocker, J.L., Konstan, J.A., Terveen, L.G., Riedl, J.T.: Evaluating collaborative filtering recommender systems. ACM Trans. Inf. Syst. (TOIS) **22**(1), 5–53 (2004)
65. Campos, J., Figueiredo, A.D.: The Serendipity Equations. Navy Center for Applied Research in Artificial Intelligence, Naval Research Laboratory, Washington, DC, Technical note AIC-01–003 (2001)
66. Campos, J., Figueiredo, A.D.: Programming for serendipity. In: Proceedings of the AAAI Fall Symposium on Chance Discovery – The Discovery and Management of Chance Events (2002)
67. Campos, J.M.M., De Figueiredo, A.C.D.: U.S. Patent No. 7,319,998. Washington, DC: U.S. Patent and Trademark Office (2008)
68. Erdelez, S.: Information encountering: it's more than just bumping into information. Bull. Assoc. Inf. Sci. Technol. **25**(3), 26–29 (1999)
69. Bogers, T., Björneborn, L.: Micro-serendipity: meaningful coincidences in everyday life shared on Twitter. iConference 2013, pp. 196–208 (2013)
70. Williamson, K.: Discovered by chance: the role of incidental information acquisition in an ecological model of information use. Lib. Inf. Sci. Res. **20**(1), 23–40 (1998)
71. Loepp, B., Hussein, T., Ziegler, J.: Choice-Based Preference Elicitation for Collaborative Filtering Recommender Systems (2014)
72. Nagulendra, S., Vassileva, J.: Understanding and controlling the filter bubble through interactive visualization: a user study. In: Proceedings of the 25th ACM Conference on Hypertext and Social Media, pp. 107–115. ACM (2014)
73. Nagulendra, S., Vassileva, J.: Providing awareness, explanation and control of personalized filtering in a social networking site. Inf. Syst. Front. **18**(1), 145–158 (2016)
74. Munson, S.A., Lee, S.Y., Resnick, P.: Encouraging Reading of Diverse Political Viewpoints with a Browser Widget. In: ICWSM (2013)
75. Kirkpatrick, M.: Facebook's Zuckerberg says the age of privacy is over, p. 198 (2010)

76. Verbert, K., Manouselis, N., Ochoa, X., Wolpers, M., Drachsler, H., Bosnic, I., Duval, E.: Context-aware recommender systems for learning: a survey and future challenges. IEEE Trans. Learn. Technol. **5**(4), 318–335 (2012)
77. Carr, P.L.: Serendipity in the stacks: libraries, information architecture, and the problems of accidental discovery. Coll. Res. Libr. **76**(6), 831–842 (2015)
78. Makri, S., Bhuiya, J., Carthy, J., Owusu-Bonsu, J.: Observing serendipity in digital information environments. Proc. Assoc. Inf. Sci. Technol. **52**(1), 1–10 (2015)
79. Dantonio, L., Makri, S., Blandford, A.: Coming across academic social media content serendipitously. Proc. Am. Soc. Inf. Sci. Technol. **49**(1), 1–10 (2012)
80. Asghar, H.M.: Measuring information seeking through Facebook: scale development and initial evidence of information seeking in Facebook scale (ISFS). Comput. Hum. Behav. **52**, 259–270 (2015)
81. Sun, T., Zhang, M., Mei, Q.: Unexpected relevance: an empirical study of serendipity in Retweets. In: ICWSM (2013)

Information Mining from Public Mailing Lists: A Case Study on IETF Mailing Lists

Heiko Niedermayer(✉), Nikolai Schwellnus, Daniel Raumer, Edwin Cordeiro, and Georg Carle

Chair of Network Architectures and Services, Department of Informatics, Technical University of Munich, Munich, Germany
{niedermh,schwelln,raumer,cordeiro,carle}@in.tum.de

Abstract. Public mailing lists, such as the mailing lists used by the IETF for Internet Standardization, can be used as big real world data set for analysis of social interactions. However, volatile participation and the usage of mail addresses as changeable pseudonyms constitute a challenge for data mining in these data. We conducted a case study of mailing list analysis wherein we address the consistent identification of a person with all of her contributions to be used as panel data. Based on the postings of individuals on different mailing lists, correlations between standardization areas in the IETF groups can be computed. Isolated and meshed standardization areas can be identified.

Keywords: Mailing lists · Identity deduplication · Clustering · Standardization

1 Introduction

Open mailing lists provide a vast area of open data for research on activities of the related groups. We have been studying the standardization efforts of the Internet Engineering Task Force (IETF) wherein open technical discussions lead to standardization of Internet technologies. Since these lists are open and can be accessed and joined via simple registration, a lot of different people from all over the world with different and changing backgrounds are represented. While the IETF and its contributing individuals are active for decades, the used mail addresses functioning as pseudonyms may change, e.g. due to the change of company affiliation. Thus, identifying an individual on the list simply by her e-mail address is not sufficient. Additionally, spam accounts and management accounts may also influence the data and they need to be addressed.

Our contribution is as follows: We applied and adapted the approach of Jensen et al. [6] for the deduplication of users. We analyze the outcome. Furthermore, having ensured that persons had been properly deduplicated, we present some results that can be generated from this data.

I. Kompatsiaris et al. (Eds.): INSCI 2017, LNCS 10673, pp. 301–309, 2017.
https://doi.org/10.1007/978-3-319-70284-1_23

2 Related Work

At the 3rd International Conference on Internet Science, we presented our initial dataset of the publicly available information of the Internet Engineering Task Force (IETF) [10]. We presented a first analysis on this dataset and showed the influence of external occurrences like the Snowden leaks on security related standardization activities and vice versa the influence of internal IETF activities on the outside world in the social media service Twitter. This work is intended to update on the grown datasets and extended analysis.

Bettenburg et al. [1] used off-the-shelf algorithms to analyze data from mailing lists. They did not find any off-the-shelf solution for clustering of multiple identities by a single person but stated that sibling identities render social analysis useless. If not solved, persons appear multiple times in the data and her actions may be considered independent actions from different individuals.

The basic ideas for deduplication of persons on mailing list data is given by Bird et al. [2]. We based our approach on the work from Jensen et al. [6]. Maijuan et al. [9] present a metric to select the best name alias out of several grouped mail identities.

Jensen et al. [6] study the behaviour of new users (newbies) on a set of open source mailing lists. They found that these newbies got replies quickly. The replies are helpful most of the time and only a small percentage of rude replies occurs. Junior et al. [7] use neurolinguistic theory to obtain and mine information about developers on a mailing list. Chen et al. [4] generate social graphs on the basis of posts and replies on mailing lists. Toral Mar et al. [8] performed a factor analysis on an open source software mailing list. Other work centers around the case that e-mail data of an organisation is given and they try to infer the social network and status of individuals on the list, e.g. [11]. Due to different foci of the mailing lists – e.g. between software development focused and standardization focused lists, results are expected to differ.

3 The IETF Dataset

We already presented our initial dataset in previous work [10]. The current and extended database schema is shown in Fig. 1. In the following, we focus on mailing lists. The data is generated from the parsed mailing list archives of all IETF mailing lists and the published drafts and RFC documents. We extract meta information for each mail. The archives include discussions and announcements as well as spam mails. We filter out spam mails on the basis of the provided spam scores. The database stores around 2 million mails posted on 984 mailing lists from around 20,000 mail addresses, which belong to 13,439 actual persons after deduplication as explain in the paper.

To understand inter-relations between different lists, we introduce unique persons that are identified by a surrogate id. Otherwise it is not possible to recognize when someone has sent mail from different accounts. In the following, we focus on lists where standardization-relevant technical discussions occur. We exclude management lists and meeting participation mailing lists from our analysis.

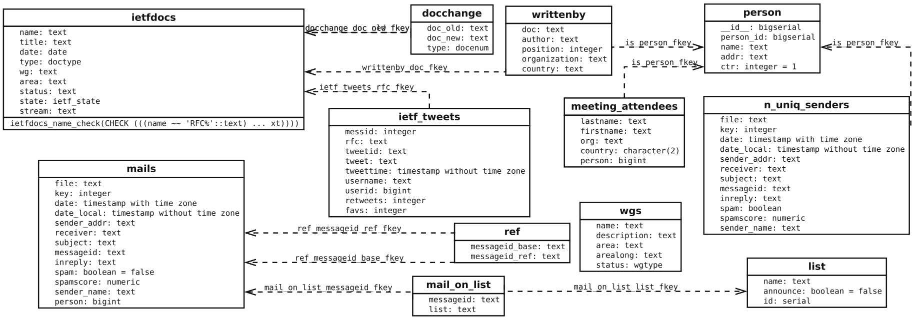

Fig. 1. Database schema

4 Identification and Deduplication of Individuals

Beside changes of mail addresses due to a change in an individual's company affiliation, we found that authors use different mail addresses on mailing lists than in the RFC documents. Other authors changed mailing addresses at certain times. These surrogate identities were not directly mappable to real individuals and presented a hurdle for further analysis. Similar problems have been reported in other cases of mailing list analysis. We follow the approach of Jensen et al. [6] with slight modifications.

4.1 Algorithm

When we fill our database with entries for a new person, they have a name given from the e-mail header and an e-mail address. If these values differ from previous entries, it will naturally be a new entry. Then the deduplication algorithm tries to find non-matching previous entries that are similar enough to merge new and old entries into one. In the database they will both exist, but share the same person identifier.

In the preprocessing, we have to normalize the name representation. This means to remove whitespace, dots, and titles in the name. Like related work we split the name into first name, middle name, and last name. Additional normalization would include normalization of special characters not found in English. We also blacklist spam names and spam e-mail addresses. We also want to exclude management accounts which might get merged due to similar names and might not contribute to the standardization discussions that we would want to analyze in our subsequent work.

We merge two persons if:

- their e-mail addresses are identical, even if names differ.
- first and last name are identical or in reversed order.
- first and last name are contained in the other e-mail address or name.

- the full names are similar enough. We use the Levenshtein edit distance to anticipate minor changes in the writing of names. The decision is then based on the metric $S = 1 - \frac{Distance(name1,name2)}{max(len(name1),len(name2))}$ and we merge if $S > 0.85$.
- if the e-mail address before @ is equal, but shorter than 6 characters, yet the full name similarity $S > 0.75$.

In addition to name and e-mail address similarity, we propose to use additional external input. Most promising for mailing lists are PGP keys and their signatures provided by keyservers given enough security-aware persons are participating in the mailing lists.

4.2 Using Information from PGP Keyservers

PGP key servers provide a ground truth for clustering e-mail addresses together as the owners of the e-mail address put them together in a PGP signature associating all of them with their PGP key. This information shows that two mail addresses belong to the same person if they share the same PGP key. It does not show, however, that two e-mail addresses are not from the same person. We can use this to improve our algorithm and make better decisions when PGP is present. If we have two entries, but both e-mail addresses use the same PGP key, we will merge the two persons into one person.

PGP entries also include the name. So, the PGP entries may give the most preferable name for the person. In our study, for 55.1% of our person entries we could find PGP keys via public key servers.

5 Persons and Groups

5.1 Statistics About the Deduplication

Overall, there are around 13,000 persons in the dataset after the deduplication. About 10,000 did not need deduplication, around 3,000 are the result of merging persons that the deduplication considered to be the same person. In normal cases the number of persons merged into one is at most slightly above 10. However, there are some larger cases in our data set. The largest one with 86 persons is from merging spam accounts that made it through the spam defense. Another issue is that persons with long identical first names might generate a situation where they get falsely merged. 46 persons got merged from more than 10 persons each, 264 from more than 5 persons each. Considering the overall number of over 13,000 persons this is a small number below 2%, but there is an increased chance among these 264 that they may have been generated from a false merge.

In our subsequent studies, in particular, the spam accounts are blacklisted and will not affect the outcome. Merging two low-profile persons will have little impact, merging a high-profile and a low-profile one as well.

5.2 In How Many Groups Do People Post?

Figure 2 shows the histogram for 2016. Most people have not been active in 2016 since our data set covers mails from as early as the year 2000. Participation has changed over time. 925 posted in only one list, 233 in two lists, 111 in three lists, and 65 in four lists. The maximum is 55 lists.

One may wonder if one can really speak of participation in a list if only few messages over a whole year were sent. If we only count lists where a person has at least 10 posts, the situation is as follows: only 368 people make it in at least one list, 90 in two lists, 25 in three lists, and 12 in four lists. The maximum is 11 lists.

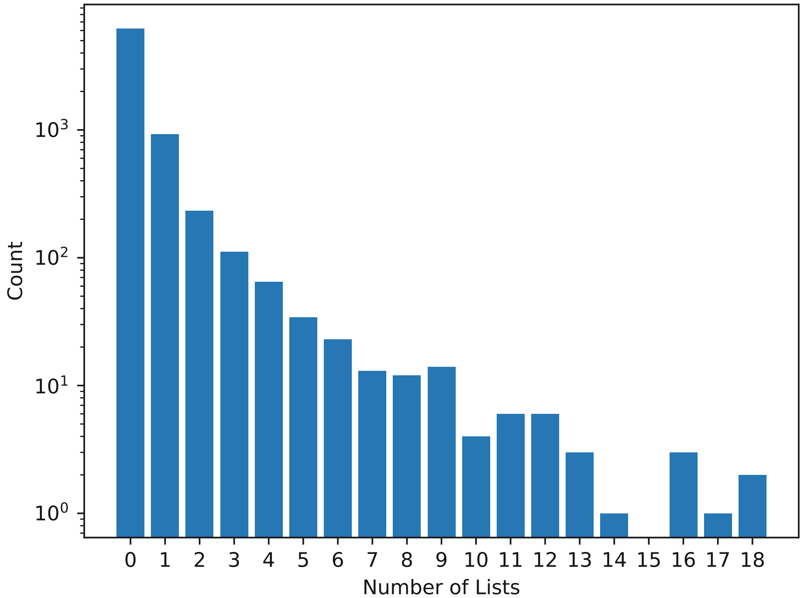

Fig. 2. Number of mailing lists a person posted in 2016

5.3 Who is Posting in Many Different Groups?

While most people only post in one or two mailing lists, there are a few individuals that post in many lists. In 2016, an individual posted in 55 lists (32 with more than 3 message), one in 39 (25 with three messages, 11 with at least 10 messages), and one in 28 (19 with at least 3 messages). All of them had occupied area director positions at the time and were responsible for a larger set of working groups which each has at least one list. The numbers drop significantly when we only count for lists with a certain amount of messages. This shows that although these persons post on a broad range of lists, their predominant activity in discussion only focuses on a subset of these lists.

5.4 How Stable is the Individual Posting Behavior over Time?

A person can either post in a group or not post in a group. We are now interested in the change from one year to another. We used a subset of our whole data that covers the last 10 years from 2007 to 2016.

In Table 1, each row refers to the number of posts in the year before. The first row means that at least one message was sent in the group by the person. The table gives the percentage of such users that one year later still post at least one message in that group (first column), 11 or more messages in that group, and so forth. In the next row, the user has posted 11 or more messages in the previous year and so forth.

Those who only posted at least one message in a group will in 40.3% of the cases post atleast one message in the following year. A majority will not. However, those who post regularly, will also post a lot of messages in the subsequent years. Even for the ones with a lot of messages (over 50) in a group, 7% will not post in the next year in this group. The likelihood that one will continue to post with a similar high rate, however, is below 50%.

Table 1. Mails $P[x_t \geq x_1 | x_{t+1} \geq x_2]$

x_1	x_2					
	1	11	21	31	41	51
1	40.3%	12.2%	7.15%	4.83%	3.42%	2.52%
11	80.4%	48.3%	33.3%	24.0%	17.6%	13.2%
21	87.5%	63.5%	48.5%	37.6%	28.5%	22.1%
31	90.1%	70.9%	58.5%	48.0%	37.9%	30.1%
41	92.4%	76.9%	66.9%	57.5%	47.8%	38.7%
51	92.9%	80.7%	71.9%	63.5%	54.6%	45.4%

5.5 How Stable is the Posting Behavior in Number of Groups?

Table 2 shows statistics for the number of groups a person posts messages in. Those who posted in at least one group also post in at least one group in the

Table 2. Lists $P[y_t \geq y_1 | y_{t+1} \geq y_2]$

y_1	y_2					
	1	2	3	4	6	11
1	58.8%	33.7%	22.6%	15.8%	8.96%	2.95%
2	74.7%	58.0%	42.8%	31.5%	19.0%	6.45%
3	84.3%	73.3%	60.1%	47.6%	30.8%	10.9%
4	89.7%	82.2%	72.6%	61.9%	43.1%	16.6%
6	94.5%	91.3%	87.1%	80.3%	63.3%	28.9%
11	96.5%	94.5%	93.9%	92.3%	86.6%	62.4%

subsequent year with 58.8%. The percentage here is higher than in the previous section because people may switch to another group and, thus, not return to a group, but still return to another group. So, taking part in the IETF overall is more stable than taking part in an individual group.

6 Inter-group Analysis

After clustering mail addresses of the same individuals, we demonstrate how the dataset can be used for the analysis of group relations. Therefore, we build a social graph which is a common way to analyze human interaction (cf. [3]). We define two groups as related when individuals posting in one group are also posting in the other group. Based on this relationship, we build a directed graph wherein each mailing list represents a node. The edge weights are computed as one minus the reciprocal of the sum of mails on the other mailing list from each poster on the origin mailing list in the fifth potency.

Figure 3 shows a spring force graph representation plotted with the Python NetworkX library. The thickness of the graph at the origin side of the edge represents the edge weight. We filtered working groups with less than 100 posts in 2016. And only added edges with an edge weight of at least 0.1.

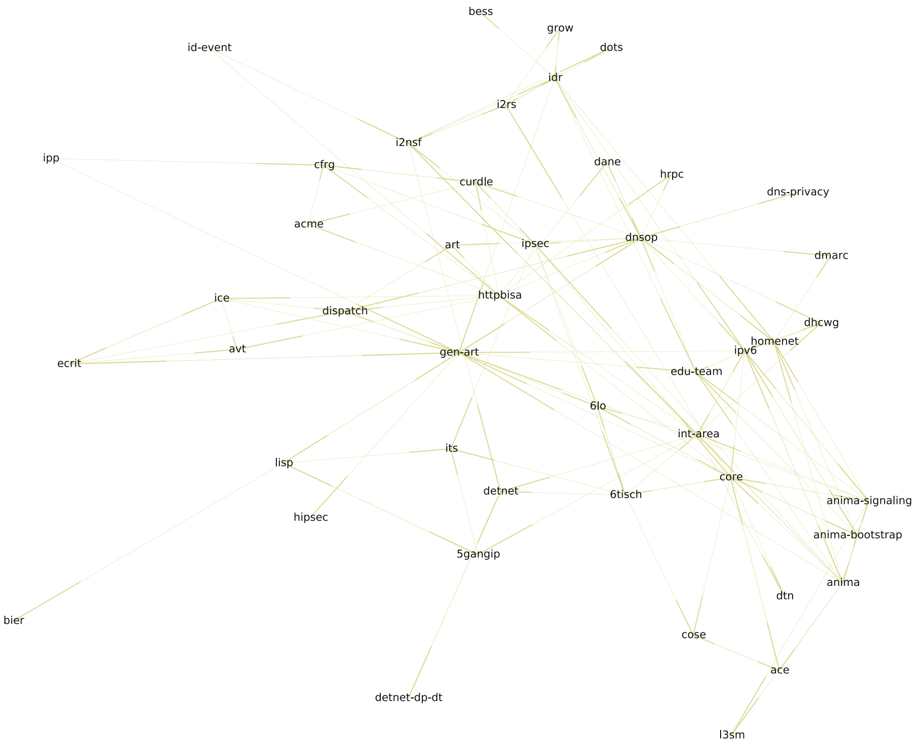

Fig. 3. Spring force graph of working groups

Although not relying on domain specific knowledge, the graph shows relations that are expected due to the relation of topic of the working groups or the IETF internal structure. The General Area Review Team (Gen-ART) mailing list takes a central position. On this list, leaders of working groups discuss and review documents with an IETF wide perspective. Also some non-general area related working groups take central positions like the Domain Name System Operations (dnsop) or Constrained RESTful Environments (core) working group. This reflects the strategical importance of these activities for other working groups.

The graph also shows subgroups with stronger intra-group links. Mailing lists like anima-signaling, anima-bootstrap and anima (all concerned with facets of autonomic networking) are positioned close to each other. Another subgroup is constituted by the security related working groups dane, ipsec, dnsop, and httpbisa; a third example is i2nsf, i2rs, and idr which are concerned with global routing. The graph also allows to identify weak ties. In analysis of social graphs, weak ties received high attention [5] due to their potential high effects that they can have on the groups connected via them. The groups bier, bess, hipsec, detnet-dp-dt, and dns-privacy only have one link to the other groups. This supposes that the connection has a strong character denoting high influence.

A closer look at the edges reveals another aspect of the relations: only some of them are symmetric (e.g. dnsop and dns-privacy). The relation between ipv6 and dnsop is asymmetric: while ipv6 members are also active in discussions in dnsop this statement cannot be made for members of dnsop that are active in ipv6.

7 Conclusions

We have presented the mailing list aspects of our IETF data set. We motivated the need for deduplication of persons and presented an adapted approach used in related work. We propose to enhance the mechanism using external sources like PGP keys from public keyservers. In a tech-savvy group like the IETF chances are high that such keys exist and improve the data set. Subsequently, we presented some statistics on how persons contribute to individual lists and to multiple lists. Finally, we provide a first visual inter-group analysis on the basis of overlap of group participants.

Indeed, contributors to IETF mailing lists are tech-savvy people. PGP keys are less likely to exist in less technical areas. Our data set covers a long period of time. Thus, changes in affiliation and email address are likely to occur. Furthermore, the users are spread around the world and issues with non-English names play a significant role. This specificity of the studied mailing lists surely has to be considered when transfering our results to other areas.

References

1. Bettenburg, N., Shihab, E., Hassan, A.E.: An empirical study on the risks of using off-the-shelf techniques for processing mailing list data. In: 2009 IEEE International Conference on Software Maintenance, pp. 539–542. Edmonton, Alberta, Canada, September 2009

2. Bird, C., Gourley, A., Devanbu, P.T., Gertz, M., Swaminathan, A.: Mining email social networks. In: Proceedings of the 2006 International Workshop on Mining Software Repositories, MSR 2006, pp. 137–143. Shanghai, China, May 2006
3. Borgatti, S.P., Mehra, A., Brass, D.J., Labianca, G.: Network analysis in the social sciences. Science **323**(5916), 892–895 (2009)
4. Chen, H., Shen, H., Xiong, J., Tan, S., Cheng, X.: Social network structure behind the mailing lists: ICT-IIIS at TREC 2006 expert finding track. In: Proceedings of the Fifteenth Text REtrival Conference, TREC 2006, Gaithersburg, Maryland, USA, November 2006
5. Granovetter, M.S.: The strength of weak ties. Am. J. Sociol. **78**(6), 1360–1380 (1973)
6. Jensen, C., King, S., Kuechler, V.: Joining free/open source software communities: an analysis of newbies' first interactions on project mailing lists. In: 2011 44th Hawaii International Conference on System Sciences, pp. 1–10, January 2011
7. Junior, M.C., Mendonca, M., Farias, M., Henrique, P.: OSS developers context-specific preferred representational systems: a initial neurolinguistic text analysis of the apache mailing list. In: 7th IEEE Working Conference on Mining Software Repositories, MSR 2010, pp. 126–129, May 2010
8. Marín, S.L.T., Martínez-Torres, M.R., Barrero, F.: Modelling mailing list behaviour in open source projects: the case of ARM embedded linux. J. UCS **15**(3), 648–664 (2009)
9. Meijuan, Y., Qingxian, W., Shuming, C., Xiaonan, L., Xiangyang, L.: Ranking the authority of name aliases for email users. In: 2011 Third International Conference on Multimedia Information Networking and Security, pp. 425–430, November 2011
10. Niedermayer, H., Raumer, D., Schwellnus, N., Cordeiro, E., Carle, G.: An analysis of IETF activities using mailing lists and social media. In: Bagnoli, F., Satsiou, A., Stavrakakis, I., Nesi, P., Pacini, G., Welp, Y., Tiropanis, T., DiFranzo, D. (eds.) INSCI 2016. LNCS, vol. 9934, pp. 218–230. Springer, Cham (2016). doi:10.1007/978-3-319-45982-0_19
11. Yoo, S., Yang, Y., Lin, F., Moon, I.C.: Mining social networks for personalized email prioritization. In: Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD 2009, New York, pp. 967–976. ACM, June 2009

Evaluation of Linked, Open Data Sources for Mining Adverse Drug Reaction Signals

Pantelis Natsiavas[1,2(✉)], Nicos Maglaveras[1,2], and Vassilis Koutkias[2]

[1] Lab of Computing, Medical Informatics and Biomedical Imaging Technologies, Department of Medicine, Aristotle University of Thessaloniki, Thessaloniki, Greece
{pnatsiavas,nicmag}@med.auth.gr

[2] Institute of Applied Biosciences, Centre for Research & Technology Hellas, Thermi, Thessaloniki, Greece
vkoutkias@certh.gr

Abstract. Linked Data is an emerging paradigm of publishing data in the Internet, accompanied with semantic annotations in a machine understandable fashion. The Internet provides vast data, useful in identifying Public Health trends, e.g. concerning the use of drugs, or the spread of diseases. Current practice of exploiting such data includes their combination from different sources, in order to reinforce their exploitation potential, based on unstructured data management practices and the Linked Data paradigm. In this paper, we present the design, the challenges and an evaluation of a Linked Data model to be used in the context of a platform exploiting social media and bibliographic data sources (namely, Twitter and PubMed), focusing on the application of Adverse Drug Reaction (ADR) signal identification. More specifically, we present the challenges of exploiting Bio2RDF as a Linked Open Data source in this respect, focusing on collecting, updating and normalizing data with the ultimate goal of identifying ADR signals, and evaluate the presented model against three reference evaluation datasets.

Keywords: Linked Open Data · Bio2RDF · Adverse Drug Reactions

1 Introduction

Pharmacovigilance is defined as "the science and activities related with the detection, assessment, understanding and prevention of adverse effects or any other possible drug-related problems" [1]. The importance of pharmacovigilance is evident as recent evidence depicts the huge financial and public health costs of Adverse Drug Reactions (ADR) worldwide [2, 3]. The Council for International Organizations of Medical Sciences (CIOMS) VIII Working Group defines pharmacovigilance signals as "information that arises from one or multiple sources (including observations and experiments), which suggests a new potentially causal association, or a new aspect of a known association, between an intervention and an event or set of related events, either adverse or beneficial, that is judged to be of sufficient likelihood to justify verificatory action" [4]. In this context, identifying new ADR signals could be significantly

I. Kompatsiaris et al. (Eds.): INSCI 2017, LNCS 10673, pp. 310–328, 2017.
https://doi.org/10.1007/978-3-319-70284-1_24

facilitated through exploiting emerging Internet technologies, like the analysis of unstructured data sources and Semantic Web [5] technologies.

Unstructured data sources, like social media and bibliography, can be exploited for eHealth using text-mining techniques [6], as they provide vast data with high geographic coverage, time evolution and user interaction characteristics, which can be useful in identifying Public Health trends. Typically, unstructured data management technologies (e.g. Apache UIMA [7]) are used to extract useful information in an exploitable data formalism. This information is then combined with knowledge obtained from other sources for further processing and comparative analysis.

The Linked Data paradigm refers to a group of technologies and standards which facilitate the interconnection and the joint exploitation of data [8]. Therefore, using Linked Data as the data representation formalism enables the creation of a common model interconnecting a variety of heterogeneous, raw data sources. In the context of pharmacovigilance, such a data model could be used to analyze information from unstructured data sources like social media, bibliography and clinical narratives, along with genetic or biochemical information and unify them in an interlinked data processing realm. The advantages of using Linked Data can be summarized as follows:

- *Interoperability:* Linked Data are based on the Resource Description Framework (RDF)[1], which uses Uniform Resource Identifiers (URIs) to unambiguously identify resources. Moreover, RDF defines syntactic and light semantic rules, which significantly facilitate interoperable communication among various IT systems.
- *Semantics and reasoning:* Linked Data standards such as RDF, RDF Schema and Web Ontology Language[2] (OWL) define high-level semantic relationships among various types of resources (e.g. hierarchies among classes, data and object properties, cardinality restrictions on object properties, functional properties, hierarchies of properties, etc.). These relationships are based on robust mathematical background (e.g. OWL semantics are based on Description Logics) and therefore, can be used for automatic inferencing and reasoning via off-the-shelf software modules, the so-called "reasoners" (e.g. Hermit[3]).
- *Knowledge reuse:* The interconnection of the variety of publicly available data sources through Linked Data standards can significantly facilitate their reuse, further exploitation and possible extension.

Typically, a Linked Data model consists of two distinct parts. On the one hand, the *Terminology Box* (a.k.a. TBox) depicts the concepts and their relations defining each data model's semantic schema and could be characterized as the "ontology part" of the model. On the other hand, the *Assertion Box* (a.k.a. ABox) contains the instances of the TBox defined concepts with their values and can be considered as the model's "data part".

In this paper, we present the challenges and a first evaluation of building a Linked Data model to be used for Public Health data analytics. The Linked Data model presented is planned to be integrated in a platform, currently under development,

[1] https://www.w3.org/RDF/.

[2] https://www.w3.org/OWL/.

[3] http://www.hermit-reasoner.com/.

aiming to extract information from open unstructured data sources, focusing on social media and bibliographic sources [9], currently exploiting Twitter and PubMed. We also focus on the semantic normalization process among the raw data sources that the presented model has been built upon. As shown in Fig. 1, the platform applies Natural Language Processing (NLP) through Apache UIMA framework on the collected raw data and represents the results using W3C's Web Annotation Data Model Linked Data formalism to facilitate their further processing through Semantic Web technologies.

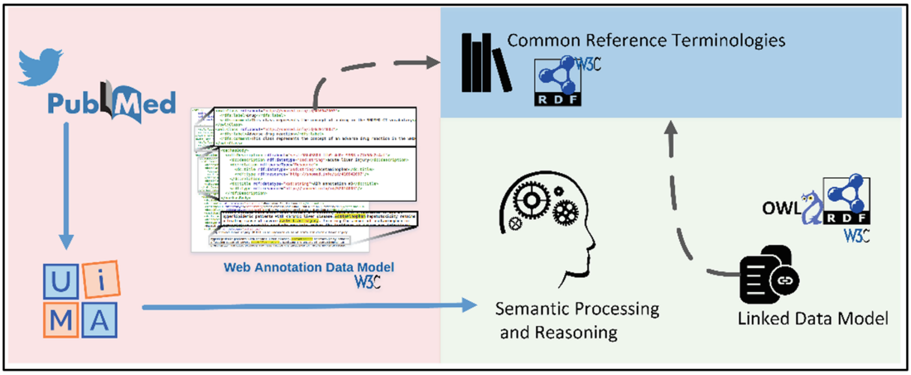

Fig. 1. Use of the presented Linked Data model in the context of unstructured data analysis platform. Dashed arrows refer to interlinking data through RDF references and normal arrows depict the flow of information.

The paper is structured as follows: In Sect. 2, we refer to related work, focusing on building and exploiting Linked Data models via Semantic Web technologies in the life sciences with emphasis on drug-related applications. In Sect. 3, we present the methodology applied regarding the definition of the proposed model and its main design goals, as well as the methodology for its evaluation. In Sect. 4, we present the results concerning the evaluation of the proposed model against its main design goals. Finally, in Sect. 5, we conclude the paper and discuss challenges, additional usage scenarios of the presented data model and future work.

2 Related Work

2.1 Linked Data Models in Life Sciences

Linked Open Drug Data (LODD) [10] has been one of the first major efforts to create an open Linked Data repository for drug research. It has been hosted as a project under the W3C's Semantic Web for Health Care and Life Sciences Interest Group (HCLS IG) and partly focused on using state-of-the-art semantic link discovery techniques for interlinking the published datasets[4].

[4] https://www.w3.org/wiki/HCLSIG/LODD/Interlinking.

Bio2RDF [11] is one of the most well-known and widely used open Linked Data sources in life sciences. It has been used in several projects, as it includes a variety of data sources transformed to RDF through a well-defined transformation process. These data sources include pharmacology-oriented data sources among others, such as DrugBank [12], FDA Structured Product Labels [13], data from ClinicalTrials.gov [14], SIDER [15] etc. The Bio2RDF data are generated from the original data sources, regardless of the initial data representation formalism, using conversion scripts following explicitly defined conventions[5]. This conversion process does not focus on enhancing the produced data model's semantics (e.g. add hierarchies of concepts describing the produced data semantics), but rather on converting raw data to the RDF formalism.

One of the most prominent European projects exploiting Linked Data technologies for drug-oriented applications is OpenPHACTS[6]. OpenPHACTS builds an "Open Pharmacological Data Space" to facilitate the discovery of new drugs. It collects data from several heterogeneous sources focusing on biochemical data [16]. The ultimate goal is to create an open platform, namely, the OpenPHACTS Discovery Platform [17], integrating data from multiple sources, accessible through an API [18]. It should be noted that OpenPHACTS employs data sources that are originally in RDF format and does not deal with converting data from other available formalisms.

DINTO [19] is a recently developed Linked Data model integrating several drug-oriented data sources (namely DrugBank, ChEBI and SIDER) in an OWL data model. DINTO has been used in combination with Semantic Web Rule Language (SWRL)[7] to facilitate drug-to-drug interactions (DDI) discovery.

Wikidata is a general purpose Linked Data repository, not restricted on Life Sciences oriented data. However, Wikidata contains a large dataset of gene annotations and health-oriented applications are one of its key application domains [20, 21].

Finally, the LinkedDrugs dataset [22] has been recently developed integrating drug product data using trade names and exploiting data from national regulatory authorities from 23 countries. While the main use case of LinkedDrugs concerns marketing, it could be also exploited for other use cases, as its data refer to DrugBank and use the Anatomical Therapeutic Chemical (ATC) classification system and, therefore, it can be interlinked with other datasets.

2.2 Application of Linked Data Paradigm in Pharmacological Research

The idea of linking heterogeneous datasets using technologies like RDF to support drug research has been presented by Boyce et al. [23]. As a follow-up, researchers from this team presented an openly accessible endpoint of searchable Linked Data focusing on pharmacovigilance and linking them with clinical evidence exploiting the Observational Medical Outcomes Partnership (OMOP) Common Data Model, as part of the Observational Health Data Sciences and Informatics (OHDSI) initiative [24]. Furthermore, the PredicTox project has presented a Linked Data approach exploiting

[5] https://github.com/bio2rdf/bio2rdf-scripts/wiki/RDFization-Guide.

[6] https://www.openphacts.org/.

[7] https://www.w3.org/Submission/SWRL/.

ontologies and terminologies to combine various data sources for ADR prediction [25], while Zhu et al. explored Semantic Web approaches on top of an OWL model of PharmGKB, aiming at repositioning breast cancer drugs [26].

Several Linked Data models have been also employed by online services. In particular, Kozac et al. presented an online service called Drug Encyclopedia, providing physicians the ability to search and browse clinically relevant information about medicinal products and drugs [27]. Drug Encyclopedia is based on a Linked Dataset relying on DrugBank, ATC and the Medical Dictionary for Regulatory Activities (MedDRA) among other knowledge bases. Doulaverakis et al. presented GalenOWL, a web service able to answer queries regarding DDIs based on Apache Jena, Sesame and OWLIM Lite to perform reasoning [28]. The same team, has also presented Panacea [29], an online recommendation system, facilitating drug-drug and drug-diseases interaction discovery.

Furthermore, several papers have focused on combining Linked Data with advanced data processing techniques. Dalleau et al. assembled a Linked Dataset using data relative to pharmacogenomics and applied Machine Learning techniques aiming at suggesting valid genes that may be involved in drug response variability [30]. Zong et al. presented a Machine Learning approach to mine drug-target associations in a Linked Data model employing DrugBank among others [31]. Focusing also on drug-target associations, in [32] a statistical model exploiting a Linked Data model is presented, employing a graph theory oriented approach. Similarly, Muñoz et al. employed SIDER and DrugBank, among other data sources, and exploited graph similarity metrics, in order to identify ADRs among drugs with similar behavior [33].

3 Methods

In this section, we present the main design goals of the proposed Linked Data model and how the methodology of building and evaluating it corresponds to them. Furthermore, we present the details of the evaluation process designed to confirm the richness of the produced Linked Data model and its prospects regarding ADR signal confirmation.

3.1 Model Design

As depicted in Fig. 1, the presented data model's primary use case concerns facilitating the linkage of the platform's results (extracted from unstructured data sources through the NLP process) with other well-known data sources and enabling semantic reasoning upon them, focusing on drug-related Public Health applications. To this end, the main desired features of the proposed model are:

- *Data richness:* The data model must be able to accommodate a vast amount of data, in order to be useful in real-world operations. Therefore, the open data sources included should add a rich ABox in the overall Linked Data model.
- *Semantic richness:* Well-defined semantics defined in Linked Data formalisms (namely, RDF and OWL) is one of the paradigm's major advantages, as they are the basis of the available automatic inference capabilities. Therefore, the designed data model should provide a rich TBox based on which automatic inferencing would be feasible and could add value.

- *Semantic Normalization:* Exploiting heterogeneous data sources through a common semantic infrastructure, requires their normalization through interlinking of their semantically matching data instances. Various data sources could refer to the same concept or object (e.g. a drug), using different identifiers, encodings, names, etc., while normalizing these data is an inherent part of the linking process.
- *Scalability:* The data model should be able to scale on two dimensions: (a) adding new data sources that could provide data or semantic enrichment, and (b) being computationally effective, meaning that it can allow the execution of the required processing in reasonable time.
- *Up-to-date maintenance:* The data model should be frequently updated, in order to accommodate potential changes of the raw data sources. As the various interlinked data sources are improved, enriched, validated, etc., the respective changes should also reflect on the produced overall data model.

Typically, in order to ensure the quality of a designed ontology, knowledge engineering methodologies are applied [34]. The presented Linked Data model is designed and implemented following the NeOn methodology. NeOn is a "scenario-based methodology that supports different aspects of the ontology development process, as well as the reuse and dynamic evolution of networked ontologies in distributed environments, where knowledge is introduced by different people (domain experts, ontology practitioners) at different stages of the ontology development process" [35].

Several data sources are employed in the current version of the model, emphasizing on the data sources available from Bio2RDF as they are already parsed in RDF format and can be directly imported in the storage infrastructure. More specifically, DrugBank, SIDER, Linked SPL, PharmGKB [36] and ClinicalTrials.gov data have been directly imported from Bio2RDF, while STITCH [37] has been converted to RDF manually. STITCH, PharmGKB and ATC [38] have been used so far only to semantically interlink the data in the rest of the data sources, while the rest of the data sources have been also validated against pharmacovigilance scenarios. The finally produced data model contains millions of RDF triples providing a vast amount of data to be used in the target scenarios. We have evaluated the data richness of the data model through a custom evaluation process (described in detail in the section entitled "Evaluation process" below), which assesses both positive and negative ADR signals, in order to confirm that the presented data model can be exploited for pharmacovigilance.

The produced data model is semantically enriched through its interconnection with external, widely accepted knowledge bases (i.e. ontologies, thesauri, hierarchies etc.), namely, the Unified Medical Language System (UMLS) [39], ATC, RxNorm [40] etc. These knowledge bases are increasingly made available in some kind of Linked Data formalism through platforms like BioPortal [41]. It should be noted that the ATC hierarchy is included in the presented model as it also facilitates interlinking of the employed data sources.

Semantic normalization of heterogeneous data sources is a far from trivial task. Exploiting one of the key Linked Data paradigm features, the data included in the presented data model systematically refer to external, widely-accepted, common reference terminologies. Figure 2 depicts the interconnection between the data sources used in the presented data model and specific external common reference terminologies, namely, RxNorm, ATC, UMLS. STITCH and PharmGKB have been also used to interlink the data from the rest of the data sources. This design decision has significantly enhanced the interlinking between the various data sources and therefore the semantic normalization of the data provided from heterogeneous data sources. As shown in Fig. 2, most of the employed data sources reference at least one external data source and the interconnection among the data model's sources and the external common reference terminologies creates a large knowledge graph semantically interconnecting most of these datasets (ClinicalTrials.gov is the only exception) in one conceptual model.

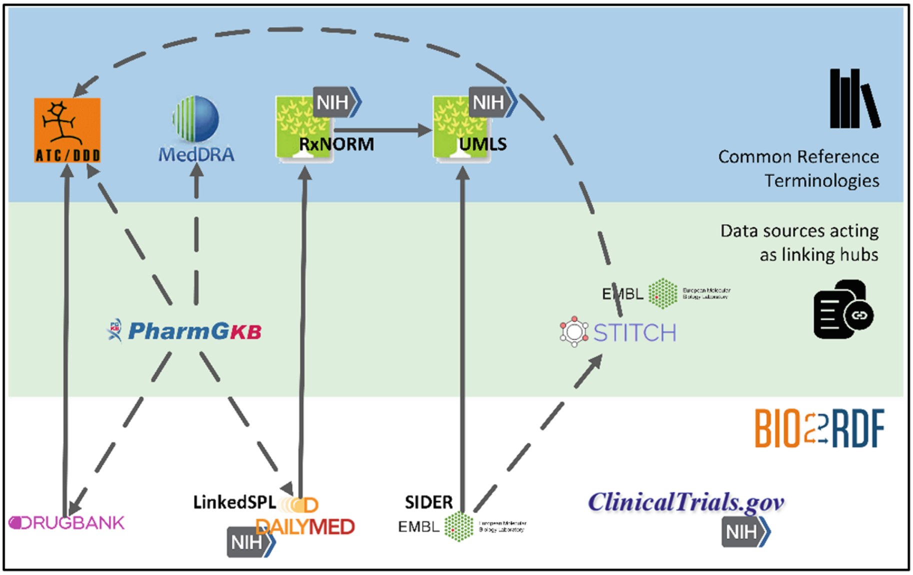

Fig. 2. Overview presentation of the data model. Dashed lines refer to interlinking paths created due to the use of PharmGKB, ATC and STITCH.

While other Linked Data sources could be employed in the produced data model, we focused on the sources that could facilitate the application of ADR signal investigation based on the textual reference of the medical condition and the drug in social media and literature. Therefore, biological and biochemical data sources which could be employed are currently considered out of the targeted application scope. For example, while data from PharmGKB and STITCH could be significant for pharmacovigilance, they could not be exploited based on textual input. Therefore, they have only been used as interlinking hubs among the rest of the data sources. It should be noted that SIDER does not directly refer to the drug identified, but only to the respective active ingredient's compound STITCH ID. Therefore, STITCH has been used to allow this interconnection between SIDER and ATC, implicitly identifying the referred drug substance.

Regarding the design decisions concerning the construction of the proposed data model and their impact on scalability, we decided to use the “Graph per Source” design pattern [42], which corresponds to building and maintaining one independent named RDF graph per each data source, as shown in Fig. 3.

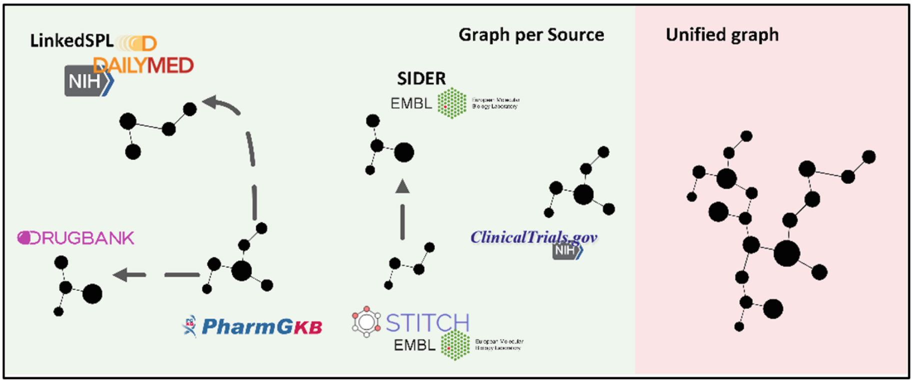

Fig. 3. Comparing “Graph per Source” design pattern (left part) with a unified monolithic named RDF graph (right part).

The “Graph per Source” design pattern improves the model’s scalability with respect to:

1. *Addition of new data sources:* Adding a new data source requires adding a new graph without engaging the already existing ones which could be used as usual.
2. *Computational effectiveness:* Organizing all the available data in one monolithic RDF graph (right part of Fig. 3) would typically result in slower retrieval times in query executions, as the size of the graph would lead to traversing a huge graph. Moreover, due to the principle of data locality [43], data from a specific data source, corresponding to one RDF graph, would be used in most of the cases for computationally intensive actions (e.g. multiple queries). Therefore, using this design pattern is expected to lead to better response times, as each query would target a smaller RDF graph, a lot quicker to traverse, in order to return the query results.
3. *Maintenance:* Finally, the “Graph per Source” design pattern facilitates the frequent synchronization of the presented data model with the raw data sources regarding up-to-date maintenance. While such a process has not yet been implemented, given that the update process of each raw data source would only refer to a specific RDF graph and not to the whole model, keeping data of each source in a separate RDF graph would significantly facilitate and accelerate the overall process. Furthermore, error recovery would also be facilitated, as a synchronization error would only be restricted to a specific named RDF graph and leave the rest of the data model intact.

3.2 Evaluation Process

The process of evaluating the data richness of the proposed data model focuses on identifying or rejecting candidate ADR signals and comparing the results against three reference evaluation datasets [44–46]. These datasets were designed to be used for the evaluation of ADR signal identification methods and contain both positive and negative controls. The use of three different evaluation datasets was decided, in order to ensure that our evaluation will be wide and diverse, covering as many drugs as possible and conditions with a sufficient number of controls, both positive or negative. In Table 1, the number of controls in each one of the evaluation datasets is presented.

Table 1. Number of controls in the reference evaluation datasets.

Dataset	Positive controls	Negative controls	Total
Harpaz et al. [44]	62	75	137
Ryan et al. [45]	165	233	398
Coloma et al. [46]	44	45	89
Totals	271	353	624

In the evaluation process, an ADR candidate signal can be described by multiple pairs of medical conditions and drug synonyms, given that medical conditions may be encoded using MedDRA, the ICD classification etc., or another textual description, while drugs can be mentioned using their trade name, their active substance, or an encoding scheme like ATC, RxNorm, etc. The evaluation algorithm accepts as textual input multiple condition-drug pairs based on the available synonyms and evaluates if the candidate signal is (a) *confirmed*, (b) *rejected*, or (c) *unconfirmed* in case that no conclusion for the specific signal can be made.

As a first step of the evaluation process, synonyms for the drug and the condition of each candidate ADR signal are retrieved. The drug synonyms are retrieved from DrugBank, while the condition synonyms are retrieved from an online service offered by BioPortal[8]. Drug synonyms vary regarding the level of granularity, since they may refer to the actual commercial drug name (e.g. Voltaren), the active chemical ingredient (e.g. Diclofenac Acid), the main pharmaceutical category (e.g. Anti-inflammatory Agents), or the chemical family of the drug's active ingredient (e.g. Non-Steroidal), with each level of synonym related with a different field of the DrugBank database. We consider that the term "Diclofenac Acid" has higher granularity level than the term "Non-Steroidal", as the first could only refer to the specific compound, while the second could refer to a whole family of drugs. Typically, we prefer to confirm or reject an ADR candidate signal based on the synonyms with the highest possible level of granularity, namely, the drug's commercial name. Therefore, the drug synonyms retrieved are semantically prioritized based on the granularity level, in order to examine the highest granularity synonyms first. If the ADR cannot be confirmed with the

[8] http://data.bioontology.org.

highest granularity level synonyms, the algorithm proceeds in examining the synonyms corresponding to lower granularity.

In the core of the evaluation process, each pair of synonyms corresponding to the given ADR candidate signal is evaluated against each ADR reference data source independently. The respective reference data sources provided by Bio2RDF (Drug-Bank, SIDER, Linked SPL and ClinicalTrials.gov) have been manually explored, in order to identify RDF properties corresponding to identifying *drugs*, medical conditions as *indication* or as an *ADR* for the respective drug. These RDF properties have been used to form the RDF paths that the evaluation algorithm queries to either confirm the given signal through the respective reference data source, or reject as the given condition refers to an *indication* of the specific drug. All the queries are formed in the SPARQL Protocol and RDF Query Language (SPARQL) [47] targeting the specific RDF paths of each reference data source, as shown in Fig. 4. First, the respective drug record is retrieved from the reference data source. Then, if the respective RDF property indicates that the candidate signal's condition is an indication of the specific drug use, the candidate signal is *rejected*. On the contrary, if the candidate signal's condition is identified in the reference data source as an ADR, the candidate signal is *confirmed*. If all the combinations of the signal's drug-condition pair signal synonyms are examined and no conclusion can be made, the candidate signal is marked as *unconfirmed*.

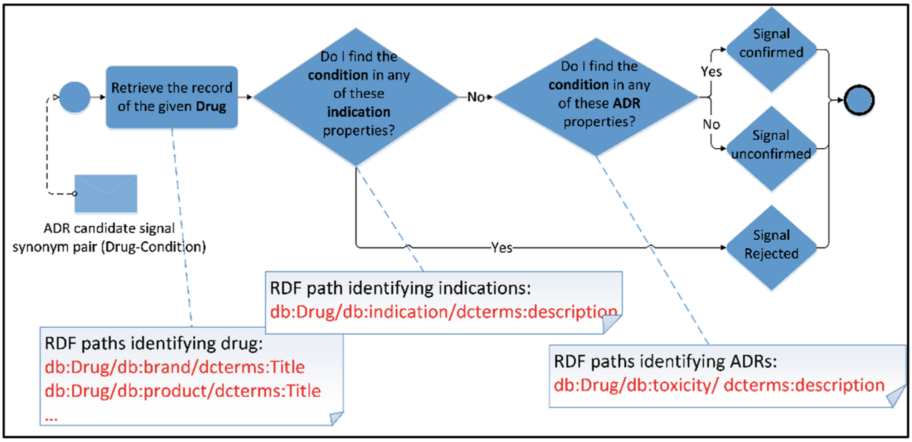

Fig. 4. Evaluation algorithm overview, depicting respective DrugBank RDF properties used in each step as an example.

It should be noted that more sophisticated candidate ADR signal evaluation algorithms could be explored. However, the main goal of the specific algorithm is to evaluate the data richness of the presented data model and not to identify or validate actual ADR signals. Therefore, the use of more advanced ADR signal identification algorithms is considered out of current work's scope.

4 Results

In this section, we present a quantified and a qualitative evaluation on most of the main design goals of the model. Furthermore, an interpretation of the quantified results is presented, referring to the detail of the original data sources.

4.1 Data Richness

Data richness in the presented Linked Data model is a top priority, as it defines whether it is actually useful for drug research applications. The evaluation of the model's data richness is performed through confirming or rejecting ADR candidate signals using the presented evaluation algorithm. The following tables present the evaluation results against the respective evaluation datasets (Tables 2, 3 and 4). For each evaluation dataset, the evaluation against each data source is presented. Furthermore, the aggregated result of the evaluation of all data sources combined is presented. It should be noted that when a signal is *unconfirmed*, it is considered as negative for the evaluation against the three evaluation datasets.

Table 2. Evaluation of data richness against Harpaz et al. [44]. SIDER presents high sensitivity, while DrugBank very high specificity.

	Aggregation	SIDER	DrugBank	LinkedSPL	ClinicalTrials.gov
Positive controls	62				
Negative controls	75				
True positive	41	48	8	30	20
True negative	20	16	68	38	48
False positive	55	59	7	37	27
False negative	21	14	54	32	42
Sensitivity	0.661	**0.774**	0.129	0.483	0.322
Specificity	0.267	0.213	**0.906**	0.506	0.64
Accuracy	0.445	0.467	0.554	0.496	0.496
F1 score	0.519	0.568	0.208	0.465	0.367

As depicted in the presented tables, it is evident that SIDER can provide results with high sensitivity (reaching about 90% for the third dataset) and DrugBank with high specificity (reaching about 97.8% for the third dataset). These results can be justified by considering the nature and the key characteristics of each raw data source:

- SIDER contains extensive information regarding ADRs, as side effects are its focus. Therefore, it is expected to be able to confirm ADRs (true positives) and lead to high sensitivity.
- DrugBank provides an extensive profile of the respective drugs, including their indications. Therefore, it is expected to be able to reject non-existing ADRs leading to high specificity.

- LinkedSPL and ClinicalTrials.gov indicate that they could facilitate the rejection of false negatives and increase specificity. However, this is not confirmed against all datasets and, therefore, it cannot be identified as a clear trend. As an overall conclusion, we can clearly confirm that the presented data model is able to be used in ADR signal identification, also confirming its data richness.

Table 3. Evaluation of data richness against Ryan et al. [45]. SIDER presents high sensitivity, while DrugBank, LinkedSPL and ClinicalTrials.gov high specificity.

	Aggregation	SIDER	DrugBank	LinkedSPL	ClinicalTrials.gov
Positive controls	165				
Negative controls	233				
True positive	45	129	29	24	38
True negative	146	95	208	169	190
False positive	87	138	25	64	43
False negative	120	36	136	141	127
Sensitivity	0.273	**0.782**	0.176	0.145	0.230
Specificity	0.627	0.408	**0.893**	**0.725**	**0.815**
Accuracy	0.480	0.563	0.595	0.485	0.573
F1 score	0.303	0.597	0.265	0.190	0.309

Table 4. Evaluation of data richness against Coloma et al. [46]. SIDER presents very high sensitivity, while DrugBank very high specificity.

	Aggregation	SIDER	DrugBank	LinkedSPL	ClinicalTrials.gov
Positive controls	44				
Negative controls	45				
True positive	31	40	13	24	18
True negative	15	5	44	20	31
False positive	30	40	1	25	14
False negative	13	4	31	20	26
Sensitivity	0.705	**0.909**	0.295	0.545	0.409
Specificity	0.333	0.111	**0.978**	0.444	0.689
Accuracy	0.517	0.506	0.640	0.494	0.551
F1 score	0.590	0.645	0.448	0.516	0.474

4.2 Semantic Richness

Semantic richness can be depicted by the number of RDF classes and properties of the model implying a semantically rich TBox. Furthermore, the links with concepts defined in external ontologies/thesauri/terminologies increases the semantic richness of the model, as it enables the exploitation of the semantics included in the externally referenced knowledge structures. Table 5 presents some quantitative results, which are indicative of the presented model's semantic richness.

Table 5. Semantic richness metrics

	Classes	Properties	External reference properties
SIDER	15	22	2
DrugBank	104	114	25
LinkedSPL	5	104	3
ClinicalTrials.gov	63	157	0
Totals	**187**	**397**	**30**

Apart from the quantitative analysis, based on the empirical use of the model we can clearly conclude that, while the produced model is characterized by a very rich ABox, its TBox could be significantly enhanced. The model could be further semantically enriched by including the concept hierarchies of external ontologies (e.g. UMLS), thus supporting automated reasoning upon them.

4.3 Semantic Normalization

Regarding the *semantic normalization and linking* efficiency of the presented model, it can be partially depicted through the number of external reference properties on Table 5, as the respective data sources can be interlinked through referencing external terminologies/ontologies/thesauri (as shown in Fig. 2).

A path for linking the data from the various data sources is by using PharmGKB as a linking hub to many knowledge sources, as depicted with red-dashed arrows in Fig. 5, since PharmGKB provides multiple linking references.

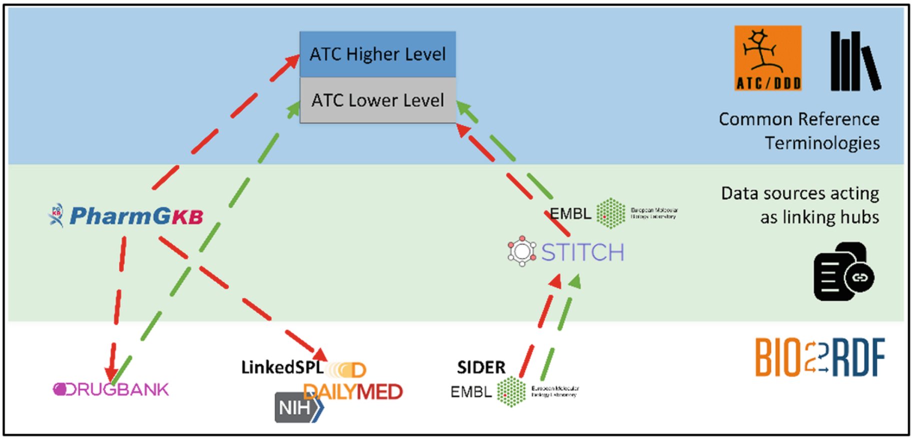

Fig. 5. Interlinking through ATC codes. Red-dashed arrows depict linking using PharmGKB as a hub and green-dashed arrows depict linking between SIDER and DrugBank using STITCH as a linking hub. (Color figure online).

Table 6. Semantic normalization and interlinking of various data sources using PharmGKB as linking hub

	Total number of drugs	PharmGKB references	Percentage
SIDER	1,593	1,593	100%
DrugBank	8,054	1,625	20.17%
LinkedSPL	51,305	886	1.72%

Table 7. Semantic normalization and interlinking of SIDER and DrugBank through ATC

	Total number of drugs	ATC references	Percentage of ATC references	Common ATC codes	Linking percentage
SIDER	1,593	1,593	100%	1049	65.85%
DrugBank	8,054	1,739	21.59%		13.02%

As shown in Table 6, PharmGKB can only partially act as an interlinking hub, normalizing content of the various data sources. The exception of SIDER is explained because PharmGKB points to the higher levels of the ATC hierarchy, referring to more abstract concepts. On the other hand, SIDER records point to the lower levels of the ATC concept hierarchy. Therefore, they can be considered matched, due to the semantic coherence implied through the ATC hierarchy, given that a concept of the lower level (more specific term) can be considered semantically included in the concepts of the above layer (more abstract term).

Another path that could be used to interlink and semantically normalize the data is the use of ATC codes. ATC codes could be used, to normalize the drug references between the two main data sources of the model, i.e. SIDER and DrugBank, as shown in Fig. 5.

The data presented in Tables 6 and 7 depict that especially in the case of SIDER, the interlinking of data reaches a significant level. However, the interlinking between the rest of the data sources can be improved.

4.4 Scalability

The evaluation process applied to verify the model's data richness, can also be used to verify its performance for data processing. The execution of the presented evaluation process corresponds to thousands of SPARQL queries to each data source and, therefore, could provide a good indicator of the model's computational efficiency. For each pair of synonyms for the drug and condition of interest, three SPARQL queries are executed (one to retrieve the drug record, one to query the drug indications and one to query respective ADRs, as shown in Fig. 4). Their execution time depends on the target RDF graph's size and the query itself which is adapted to each reference data source's RDF paths. Average execution times of the three queries for each reference data source could provide an indicative measure of the model's overall computational effectiveness. The numbers presented in Table 8 refer to the execution of the evaluation process on an Ubuntu Server Virtual Machine with 32 GBs of RAM. It should be noted that the queries are run against a Virtuoso RDF Triple Store[9] (version 07.20.3214) using Java as a programming interface to run the multiple queries in a multi-threading fashion.

Table 8. Average execution time (in seconds) of the three main queries against the data model

	1st query	2nd query	3rd query
SIDER	2.79	22.12	21.53
DrugBank	1.37	1.42	1.55
LinkedSPL	23.58	21.73	22.86
ClinicalTrials.gov	1.70	3.33	3.26

The presented execution times can only be considered indicative and could significantly improve on better infrastructure or with query optimization techniques. Overall, data presented in Table 8 depict that the presented model can provide results in reasonable execution times, despite its size and data richness.

5 Conclusions and Future Work

Given that the presented model has been developed to be used as part of an overall platform exploiting results from unstructured data analysis for Public Health applications, we argue that the presented Linked Data model would add value as it combines heterogeneous data sources, allowing their uniform data processing with Semantic Web technologies. Our evaluation which targeted ADR candidate signal identification proves that the presented data model can be valuable in such Public Health applications.

Based on the process of building and using the presented data model so far, we can conclude that it satisfies the main design goals. More specifically, our evaluation results illustrated high sensitivity on SIDER and high specificity using DrugBank as a reference data source. Therefore, it is safe to assume that interlinking the data of the two data sources (among others) we could reach better sensitivity and better specificity on the unified model.

However, there is still room for improvement, mostly regarding three key points, included in our future work plans:

1. Interlinking of the various data models can be improved using NLP techniques, in order to automatically relate free-text data with well-defined concepts in widely accepted terminologies/ontologies/thesauri. Referencing such widely-accepted terminologies would implicitly facilitate *semantic interlinking* of the various data sources of our model.

[9] https://virtuoso.openlinksw.com/rdf/.

2. The presented evaluation clearly depicts that the model could be *semantically enriched*, meaning that the model's TBox can be significantly enhanced. This is a key issue for the further semantic processing of data, as referring to external ontologies with defined hierarchies of interlinked concepts could facilitate future exploitation of Semantic Web technologies, applying automatic inferencing mechanisms.
3. *Maintenance* of the model and keeping it synchronized with the raw data sources would add significant value. Such a maintenance process would have to automatically update the model, also dealing with removing records that have been deleted in the original data source, adding new records and maintaining the semantic links of the model.

Another direction of future work involves adding alternative data sources to the presented data model. More specifically, DINTO and OpenPHACTS are the two data sources considered as the model's enhancement next step. While data richness has been demonstrated in the current version of the model, combining new data sources could significantly increase the model's use cases and, therefore, its overall impact. The scenarios investigated include drug repurposing and precision medicine applications combining genetic, biochemical and evidence-based data.

Finally, we plan to adopt quality assurance mechanisms based on ontology quality metrics [48], in order to assure that the model is well-maintained and semantically-enhanced properly.

References

1. World Health Organization: WHO collaborating centre for international drug monitoring: the importance of pharmacovigilance. World Health Organization (2002)
2. Sultana, J., Cutroneo, P., Trifirò, G.: Clinical and economic burden of adverse drug reactions. J. Pharmacol. Pharmacother. **4**, S73–S77 (2013)
3. National Safety and Quality Health Service Standards Australian Commission on Safety and Quality in Health Care: Australian Commission on Safety and Quality in Health Care (ACSQHC), Sydney, September 2011
4. Council for International Organizations of Medical Sciences (CIOMS): Practical Aspects of Signal Detection in Pharmacovigilance, Council for International Organizations of Medical Sciences. Report of CIOMS Working Group VIII. CIOMS, Geneva (2010)
5. Berners-Lee, T., Hendler, J., Lassila, O., et al.: The semantic web. Sci. Am. **284**, 28–37 (2001)
6. Harpaz, R., Callahan, A., Tamang, S., Low, Y., Odgers, D., Finlayson, S., Jung, K., LePendu, P., Shah, N.H.: Text mining for adverse drug events: the promise, challenges, and state of the art. Drug Saf. **37**, 777–790 (2014)
7. Apache UIMA - Apache UIMA. http://uima.apache.org/
8. Heath, T., Bizer, C.: Linked data: evolving the web into a global data space. Synth. Lect. Semant. Web Theory Technol. **1**, 1–136 (2011)
9. Natsiavas, P., Maglaveras, N., Koutkias, V.: A public health surveillance platform exploiting free-text sources via natural language processing and linked data: application in adverse drug reaction signal detection using PubMed and Twitter. In: Riaño, D., Lenz, R., Reichert, M. (eds.) KR4HC/ProHealth 2016. LNCS (LNAI), vol. 10096, pp. 51–67. Springer, Cham (2017). doi:10.1007/978-3-319-55014-5_4

10. Samwald, M., Jentzsch, A., Bouton, C., Kallesøe, C.S., Willighagen, E., Hajagos, J., Marshall, M.S., Prud'hommeaux, E., Hassenzadeh, O., Pichler, E., Stephens, S.: Linked open drug data for pharmaceutical research and development. J. Cheminform. **3**, 19 (2011)
11. Callahan, A., Cruz-Toledo, J., Ansell, P., Dumontier, M.: Bio2RDF release 2: improved coverage, interoperability and provenance of life science linked data. In: Cimiano, P., Corcho, O., Presutti, V., Hollink, L., Rudolph, S. (eds.) ESWC 2013. LNCS, vol. 7882, pp. 200–212. Springer, Heidelberg (2013). doi:10.1007/978-3-642-38288-8_14
12. Law, V., Knox, C., Djoumbou, Y., Jewison, T., Guo, A.C., Liu, Y., Maciejewski, A., Arndt, D., Wilson, M., Neveu, V., Tang, A., Gabriel, G., Ly, C., Adamjee, S., Dame, Z.T., Han, B., Zhou, Y., Wishart, D.S.: DrugBank 4.0: shedding new light on drug metabolism. Nucleic Acids Res. **42**, D1091–D1097 (2014)
13. Hassanzadeh, O., Zhu, Q., Freimuth, R., Boyce, R.: Extending the "web of drug identity" with knowledge extracted from United States product labels. AMIA Jt. Summits Transl. Sci. Proc. **2013**, 64–68 (2013)
14. Home - ClinicalTrials.gov. https://clinicaltrials.gov/
15. Kuhn, M., Letunic, I., Jensen, L.J., Bork, P.: The SIDER database of drugs and side effects. Nucleic Acids Res. **44**, D1075–D1079 (2016)
16. Hu, Y., Bajorath, J.: Learning from "big data": compounds and targets. Drug Discov. Today **19**, 357–360 (2014)
17. Gray, A.J.G., Groth, P., Loizou, A., Askjaer, S., Brenninkmeijer, C., Burger, K., Chichester, C., Evelo, C.T., Goble, C., Harland, L., Pettifer, S., Thompson, M., Waagmeester, A., Williams, A.J.: Applying linked data approaches to pharmacology: architectural decisions and implementation. Semant. Web **5**, 101–113 (2014)
18. Groth, P., Loizou, A., Gray, A.J.G., Goble, C., Harland, L., Pettifer, S.: API-centric linked data integration: the open PHACTS discovery platform case study. Web Semant. Sci. Serv. Agents World Wide Web **29**, 12–18 (2014)
19. Herrero-Zazo, M., Segura-Bedmar, I., Hastings, J., Martínez, P.: DINTO: using OWL ontologies and SWRL rules to infer drug-drug interactions and their mechanisms. J. Chem. Inf. Model. **55**, 1698–1707 (2015)
20. Mitraka, E., Waagmeester, A., Burgstaller-Muehlbacher, S., Schriml, L.M., Su, A.I., Good, B.M.: Wikidata: a platform for data integration and dissemination for the life sciences and beyond. bioRxiv (2015)
21. Putman, T.E., Lelong, S., Burgstaller-Muehlbacher, S., Waagmeester, A., Diesh, C., Dunn, N., Munoz-Torres, M., Stupp, G.S., Wu, C., Su, A.I., Good, B.M.: WikiGenomes: an open web application for community consumption and curation of gene annotation data in Wikidata. Database (Oxford) **2017** (2017)
22. Jovanovik, M., Trajanov, D.: Consolidating drug data on a global scale using Linked Data. J. Biomed. Semant. **8**, 3 (2017)
23. Boyce, R.D., Ryan, P.B., Norén, G.N., Schuemie, M.J., Reich, C., Duke, J., Tatonetti, N.P., Trifirò, G., Harpaz, R., Overhage, J.M., Hartzema, A.G., Khayter, M., Voss, E.A., Lambert, C.G., Huser, V., Dumontier, M.: Bridging islands of information to establish an integrated knowledge base of drugs and health outcomes of interest. Drug Saf. **37**, 557–567 (2014)
24. Knowledge Base Workgroup of the Observational Health Data Sciences and Informatics (OHDSI) Collaborative: Large-scale adverse effects related to treatment evidence standardization (LAERTES): an open scalable system for linking pharmacovigilance evidence sources with clinical data. J. Biomed. Semant. **8**, 11 (2017)
25. Zaman, S., Sarntivijai, S., Abernethy, D.: Use of biomedical ontologies for integration of biological knowledge for learning and prediction of adverse drug reactions. Gene Regul. Syst. Bio. **11** (2017)

26. Zhu, Q., Tao, C., Shen, F., Chute, C.G.: Exploring the pharmacogenomics knowledge base (PharmGKB) for repositioning breast cancer drugs by leveraging web ontology language (OWL) and cheminformatics approaches. Pac. Symp. Biocomput. 172–182 (2014)
27. Kozák, J., Nečaský, M., Pokorný, J.: Drug encyclopedia – linked data application for physicians. In: Arenas, M., Corcho, O., Simperl, E., Strohmaier, M., d'Aquin, M., Srinivas, K., Groth, P., Dumontier, M., Heflin, J., Thirunarayan, K., Staab, S. (eds.) ISWC 2015. LNCS, vol. 9367, pp. 41–56. Springer, Cham (2015). doi:10.1007/978-3-319-25010-6_3
28. Doulaverakis, C., Nikolaidis, G., Kleontas, A., Kompatsiaris, I.: GalenOWL: Ontology-based drug recommendations discovery. J. Biomed. Semant. **3**, 14 (2012)
29. Doulaverakis, C., Nikolaidis, G., Kleontas, A., Kompatsiaris, I.: Panacea, a semantic-enabled drug recommendations discovery framework. J. Biomed. Semant. **5**, 13 (2014)
30. Dalleau, K., Marzougui, Y., Da Silva, S., Ringot, P., Ndiaye, N.C., Coulet, A.: Learning from biomedical linked data to suggest valid pharmacogenes. J. Biomed. Semant. **8**, 16 (2017)
31. Zong, N., Kim, H., Ngo, V., Harismendy, O.: Deep mining heterogeneous networks of biomedical linked data to predict novel drug–target associations. Bioinformatics **33**, 2337–2344 (2017)
32. Chen, B., Ding, Y., Wild, D.J., Barabsi, A., Vidal, M.: Assessing drug target association using semantic linked data. PLoS Comput. Biol. **8**, e1002574 (2012)
33. Muñoz, E., Nováček, V., Vandenbussche, P.-Y.: Using drug similarities for discovery of possible adverse reactions. AMIA Annu. Symp. Proc. **2016**, 924–933 (2016)
34. Nazario, D.C., Dantas, M.A.R., Todesco, J.L.: Knowledge engineering: survey of methodologies, techniques and tools. IEEE Lat. Am. Trans. **12**, 1553–1559 (2014)
35. Suárez-Figueroa, M.C., Gómez-Pérez, A., Fernández-López, M.: The NeOn methodology for ontology engineering. In: Suárez-Figueroa, M., Gómez-Pérez, A., Motta, E., Gangemi, A. (eds.) Ontology Engineering in a Networked World, pp. 9–34. Springer, Heidelberg (2012). doi:10.1007/978-3-642-24794-1_2
36. Whirl-Carrillo, M., McDonagh, E.M., Hebert, J.M., Gong, L., Sangkuhl, K., Thorn, C.F., Altman, R.B., Klein, T.E.: Pharmacogenomics knowledge for personalized medicine. Clin. Pharmacol. Ther. **92**, 414–417 (2012)
37. Szklarczyk, D., Santos, A., von Mering, C., Jensen, L.J., Bork, P., Kuhn, M., von Mering, C., Jensen, L.J., Bork, P., Kuhn, M.: STITCH 5: augmenting protein–chemical interaction networks with tissue and affinity data. Nucleic Acids Res. **44**, D380–D384 (2016)
38. ATC - Structure and principles. https://www.whocc.no/atc/structure_and_principles/
39. Bodenreider, O.: The unified medical language system (UMLS): integrating biomedical terminology. Nucleic Acids Res. **32**, D267–D270 (2004)
40. RxNorm Overview. https://www.nlm.nih.gov/research/umls/rxnorm/overview.html
41. Salvadores, M., Alexander, P.R., Musen, M.A., Noy, N.F.: BioPortal as a dataset of linked biomedical ontologies and terminologies in RDF. Semant. Web **4**, 277–284 (2013)
42. Dodds, L., Davis, I.: Linked data patterns (2011). http://patterns.dataincubator.org/
43. Denning, P.J.: The locality principle. Commun. ACM **48**, 19 (2005)
44. Harpaz, R., Odgers, D., Gaskin, G., DuMouchel, W., Winnenburg, R., Bodenreider, O., Ripple, A., Szarfman, A., Sorbello, A., Horvitz, E., White, R.W., Shah, N.H.: A time-indexed reference standard of adverse drug reactions. Sci. Data. **1**, 140043 (2014)
45. Ryan, P.B., Schuemie, M.J., Welebob, E., Duke, J., Valentine, S., Hartzema, A.G.: Defining a reference set to support methodological research in drug safety. Drug Saf. **36**, 33–47 (2013)

46. Coloma, P.M., Avillach, P., Salvo, F., Schuemie, M.J., Ferrajolo, C., Pariente, A., Fourrier-Réglat, A., Molokhia, M., Patadia, V., van der Lei, J., Sturkenboom, M., Trifirò, G.: A reference standard for evaluation of methods for drug safety signal detection using electronic healthcare record databases. Drug Saf. **36**, 13–23 (2012)
47. SPARQL 1.1 Overview. https://www.w3.org/TR/2013/REC-sparql11-overview-20130321/
48. Zaveri, A., Rula, A., Maurino, A., Pietrobon, R., Lehmann, J., Auer, S.: Quality assessment for Linked Data: a survey. Semant. Web **7**, 63–93 (2015)

A Data-Driven Model for Linking Open Economic Information

M. Vafopoulos[1], A. Koukourikos[1(✉)], G. Vafeiadis[1], D. Negkas[2], I. Skaros[2], and A. Tzani[2]

[1] Institute of Informatics and Telecommunications, National Centre for Scientific Research "Demokritos", Agia Paraskevi, Greece
vafopoulos@gmail.com, kukurik@iit.demokritos.gr, vafeiadis.giorgos@gmail.com

[2] University of Piraeus, Piraeus, Greece
dimneg@gmail.com, skaros.ilias@gmail.com, tzaniaggelos@gmail.com

Abstract. While public finance data are becoming openly available as part of the broader promotion of fiscal transparency, there is little effort towards maximizing their potential value by interlinking them under a concrete framework and establishing the means to extract interesting insights. The Linked Open Economy model (LOE) aims to act as a top-level conceptualization that connects economic flows with open economic data and as an adaptable and extensible underlying model for modelling different scenarios. The paper presents the LOE model, emphasizing its theoretical foundations. Furthermore, it presents the usage of the model in realistic settings, showcasing its extensibility and its ability to address interesting questions.

Keywords: Open data · Linked Data · Semantic web · Economy · Public procurement · Prices · Circular financial flow model

1 Introduction

Eight years after the call from [1] for public investments in promoting open data-driven "computational social science", their ascertainment that "…in the leading disciplinary journals in economics, sociology, and political science, there would be minimal evidence of an emerging computational social science engaged in quantitative modelling…" seems to be still valid.

Despite the plethora of open and big data initiatives that publish an impressive number of datasets, there is not yet much evidence of related scientific research. Maybe because published datasets are of low quality and relevance to research goals or established methodologies in science have not yet adapted to the so called "data revolution" [2]. In any case, it is interesting to try to bridge the gap between the mass available data and research efforts.

Particularly in the field of Economics, only lately some scholars have been starting to discuss the need of incorporating new data in empirical research (see for instance [3, 4]. The recently coined terms "Open data business models" and "Open data economy"

I. Kompatsiaris et al. (Eds.): INSCI 2017, LNCS 10673, pp. 329–343, 2017.
https://doi.org/10.1007/978-3-319-70284-1_25

analyse the economic opportunities emerging by the increasing open data provision [5] and the economic benefits of open data [6], respectively. Research in modern disciplines such as Information systems, Business Intelligence and Finance have shown faster reflexes in addressing big data issues, mainly focusing on analytics (see for instance [7–9]).

Today, the most active efforts in open economic data are coming from Open Government Data initiatives that promote innovation, transparency and accountability (e.g. Open Government Partnership [10] and Open Ownership) and mainly refer to the fields of public budgeting, procurement, spending, Official Development Assistance, subsidies and corporate information.

Tygel et al. [11] analyse open budget data initiatives and conclude that special attention should be paid to user feedback, semantics standards and linking possibilities. There are also some research efforts on semantic modelling of budgets that are mainly focused on the available data from specific countries and regions (see for instance [12–17]). In the field of data standards, the Fiscal Data Package is developed as a simple, open technical specification for government budget and spending data.

Alvarez-Rodríguez et al. [18] review the efforts of implementing semantic technologies in the field of e-procurement. Indicatively, the list of projects includes LOTED2, Public Contracts Ontology, Methods On Linked Data for E-procurement Applying Semantics (MOLDEAS) and PPROC ontology.

Recently, the Open Contracting Partnership has developed the Open Contracting Data Standard (OCDS) [19] that sets out key documents and data that should be published at each stage of a contracting process. The Standard is backed up by a documented specification describing data fields and structures that publishers should use to increase the accessibility, usability and interoperability of their disclosures.

In the field of public spending, OpenSpending.org [15] is an open platform for government expenditure tracking created by the Open Knowledge Foundation. It offers an easy system to upload, explore and share public finance data. The PublicSpending.net project [20–22] cleans, analyses and converts to LOD public spending data from seven governments, both local and national, with total value almost 1.5 trillion euros.

With regards to corporate information, Opencorporates.com [23] is an effort of aggregating company information from different countries and jurisdictions and releasing it as open data. The opencorporates team works on creating Linked Data representations out of their databases, by mapping company metadata to certified ontologies such as the Core Business Vocabulary and linking them to other data hubs, such as DBpedia.org and Geonames.

Openownership.org has been recently established to create a global beneficial ownership register.

The aforementioned initiatives are paving the way for open data in diverse aspects of economic activity, but in practice act as e-catalogues that are fragmented into topic, place and time since they do not share common standards and methodologies. In cases where open data exist, the basic obstacle is the fact that there are not even common practices for representing the main actors (e.g. payers and payees) and the type of payments. Therefore, it is impossible to interlink the available data in meaningful ways and support services and decision-making. Surely, the cost of data discovery and collection has been substantially decreased, but getting valuable insights in public finance still demands high expertise and timely efforts; and this is a serious danger that may undermine the further development of LOD in general.

Hence, it is time to guide our focus in developing more comprehensive approaches to interconnect the stylized facts of open economic data. The potential usages of such conceptualization ranges from crowd-sourced monitoring and risk assessment of public finance to real-time integration in Business Intelligence systems for more efficient resource allocation.

For instance, subsidies to public and private organizations and the provision of aid to third countries and international organizations can be considered at the same conceptual level of public procurement because they are both money transfers (payments) through a predefined process (e.g. open call). The main difference between public procurement and subsidies or aid provisioning is that in the first case the buyer (or payer) receives direct compensation (e.g. product or service) for its payments. In the case of subsidies and aid provisioning the direct benefits stay within the beneficiary while the society or specific groups of it enjoy the indirect benefits (e.g. social inclusion, economic development).

In this context, the proposed Linked Open Economy (LOE) model aims to bridge theory-driven approaches that offer generality and scalability with the more readily applicable data-driven approaches that cover specific, realistic modelling needs.

More specifically, the LOE model addresses economic open data orchestration by providing a series of coherent and scalable conceptualizations that are based on economic theory and business practice, and extending them to reflect real-world practices and requirements.

2 The Linked Open Economy Model

2.1 Model Principles

LOE builds on top of the four-sector Circular Flow of Income (CFM) conceptualization used in basic macroeconomics (Fig. 1), where major exchanges are modelled as flows of money, goods and services between economic agents falling under four distinct sectors: (1) Households, (2) Firms, (3) Government and (4) Rest of the world. These agents cooperate over their activity in specific markets of (i) Goods and Services, (ii) Factors or Resources (iii) Financial constructs and participate in flows of economic activity, either financial (direct or indirect exchange of money) or real (exchange of goods, services or factors).

Focusing on the part of CFM that involves governmental actions, we specialize the income flow process as follows: governments form and publish budgets, partially targeting projects and works that are assigned through calls for tenders. The transfer of funds, specified via signed contracts, is realized after the completion of the projects.

The following subsection elaborates on the conceptualization of agents, markets and flows in the LOE model, discussing on the correspondences with CFM and arguing on the assumptions adopted for building the model, based on the realistic restrictions posed by the nature, scope and range of the openly available economic data (marked with red bubbles in Fig. 1).

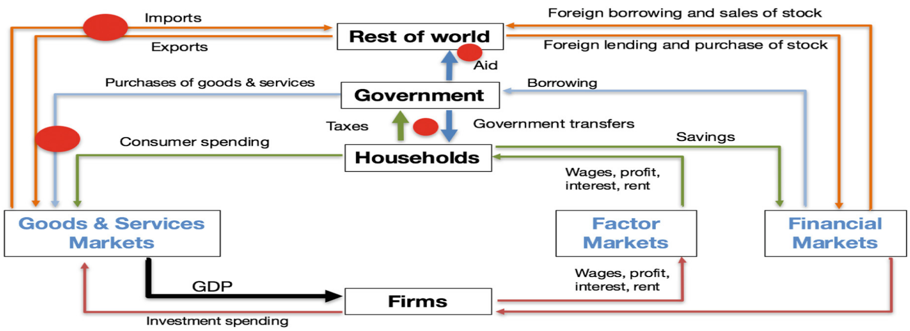

Fig. 1. Extended circular flow model

2.2 Adapting CFM to the Open Economic Data Environment

Taking into account the intricacies of the open government data, the LOE model proposes the following conceptualization for the various constructs defined in CFM.

Economic Agents

Economic agents activated in the CFM flow are represented as instantiations of a generic Agent class. Households are modelled as Persons, as the available open economic data do not include information at the level of households, but rather at the level of a single person as an economic agent. On the other hand, firms are modelled as organizations and business entities. In particular, available data involve a Business firm as a seller of services or products to the Government. Furthermore, a Business firm or a Government can benefit from Subsidies provided by a Government which, if provided to foreign countries, are considered to be International Aid. In this context, a Government acts mainly as a buyer of services and products from Business firms or - in relatively few cases - from other Governments or Governmental units. Note that, in the broader picture, a Business firm routinely acts as a buyer of goods but B2B transaction data are not publicly available.

Another limitation stemming from the scope of publicly available data is the fact that open international trade data describe solely the cumulative value of bilateral trades, with no further information on the specific Business firms that participate in import and export activities. Consequently, LOE groups the Business firms that export and import goods and services in a country as Group National Agents with the given country code, in order to capture the Rest of the World, Imports and Exports constructs in the CFM.

Flows of Economic Activity

In LOE, direct or indirect money flows are represented as Amounts exchanged between economic agents. Additionally, real flows (i.e. goods, services or factors) are modelled either as foreign Trade Activities or Public Procurement Activities. The latter activities are modelled using the following constructs, reflecting the budgetary and procurement process followed by Greek central and local government, while being straightforwardly extensible to cover other public administrations.

Expense items can be distinguished as Budget Items, Committed Items, Expense Approval Items, or Spending Items. Additionally, Revenue Items can be either Budget Items, Revenue Recognized Items, or Collected Items.

- A Budget Item models a part of a Budget, i.e. an allocation of funds for specific purposes based on an accounting system and according to an annual time plan.
- A Committed Item represents a firm, written obligation from a public organization (buyer) to provide a specified amount of funds, related to a specific Budget Item, under particular terms and conditions and for specific purposes. The recipient (beneficiary or seller) may be defined at this or at a later stage. Commitments can be issued before or after the procurement process. In the case of public budgeting in Greece, commitments are issued before the procurement process.
- The Expense Approval Item construct is introduced to model cases where further approvals are required to proceed with payments. The approvals are bureaucratic administrative decisions issued after the delivery of a contract and before a payment.
- A Spending Item represents the final stage of an Expense Budget Item. Part or the total of related expense approval items proceed to payment from buyer to seller.
- A Revenue Recognized Item represents a revenue that has been recognized by the public organization and is qualified for collection (e.g. a fine or a tax).
- A Collected Item represents the final stage of a Revenue Budget Item, i.e. part or the total of revenue budget items are collected by the public organization from a third party (e.g. taxpayers and central government).

Markets

The concept of markets is included in the LOE model primarily through open data for prices for specific products (e.g. basic goods, agricultural products and fuel), since there is no other source of publicly available data.

Taking into account the aforementioned remarks and assumptions the LOE model proposes a specialization of the CFM tailored to the specificities of the targeted domain, as depicted in Fig. 2.

The following section presents the formal specification of the LOE model, adapting the generic CFM approach to the actual data being made available.

2.3 Formal LOE Model Specification

Given the aforementioned definitions and assumptions, the LOE model formally implements the resulting conceptualization as an ontology using the Web Ontology Language (OWL) W3C specification. LOE extensively uses established ontologies and is fully conforming with the Core Vocabularies developed by the European Union.

Usage of Existing Ontologies

The following table summarizes the external ontologies used in LOE for defining the concepts and roles foreseen by the model (Table 1).

FOAF is used to describe agents responsible for specific actions as defined in LOE. Specializations relevant to LOE are foaf:Person and foaf:Organization. In the same fashion, the GoodRelations ontology is used to describe the Business Entities involved in a commercial activity, the type of their services, and the financial details of the

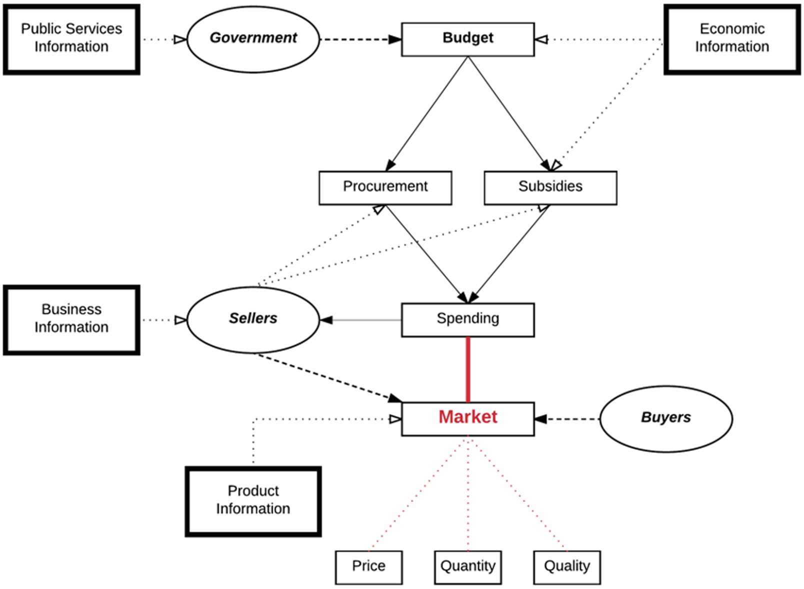

Fig. 2. CFM model specialization for public spending connected to the market

Table 1. Namespaces used at the general LOE model

Ontology	Namespace	Prefix
Friend of a friend	http://xmlns.com/foaf/0.1/	foaf
Good relations	http://purl.org/goodrelations/v1#	gr
Organizations	http://www.w3.org/ns/org#	org
Registered organization	http://www.w3.org/ns/regorg#	rov
Public contracts	http://purl.org/procurement/public-contracts#	pc
Linked economy	http://linkedeconomy.org/ontology#	elod

contract or of the payment. The Organization ontology is used to define the organizations along with their organizational units. Public Contract ontology is used to define the following types of information:

- public contracts during all stages of their existence
- procedures specifying how the details of a contract is published and how a supplier is selected
- main object of the contract (e.g. works, supplies or services)
- contract's price, depending on its stage (before or after the offer)
- award criteria that define the conditions under which the best offer will be selected and awarded, along with their weights and
- main and supplementary products or services purchased by the contract (as determined by their CPV codes).

The conceptualization of organizational entities is carried out through the usage of the Organization ontology and the Registered Organization vocabulary.

Conformance with Core Vocabularies
In order for a model to conform to the e-Government Core Vocabularies as developed by the European Union, it needs to publish a mapping to the conceptual model of the Core Vocabularies, as a self-conformance statement. The statement must cover two basic requirements:

- Each data element in the model is required to have an identifier, a label and a definition;
- The provided mapping should include the following information: Core Vocabulary identifier, Core Vocabulary version, mapping relation, Identifier, Identifier of the data element that is mapped, and a comment on the mapping.

The LOE specification includes a self-conformance statement in the form of a mapping spreadsheet, publicly available via the LOE model distribution.

3 The LOE Model in Practice

While LOE is based on a concretely defined economic abstraction as the CFM, its true aim is to bridge this theoretical foundation with the current and future reality of the open financial data ecosystem and, on an even broader context, with the Web 3.0 practices. As a testbed for assessing the model's adaptability and extensibility, it has been used as the backbone of the YourDataStories (YDS) project platform (platform.yourdatastories.eu), a framework aiming to serve stakeholders of different communities under different use cases. The following sections summarize the use cases defined in the context of YourDataStories, the respective data to be processed, and the extensions posed by the specific requirements. Detailed access information is summarized in Sect. 4.

3.1 YourDataStories Use Cases and Objectives

YDS operates over three distinct but connected use cases, all handling economic data of different scope and granularity and targeting different operations where the usage of open data can produce significant added value.

YourDataStories Pilot #1: Follow Public Money. Often, information about public projects resides in not connected systems owned by different ministries and public agencies. In the case of Greece, this information can be mainly found in the form of open data in two systems: (1) the NSRF portal (anaptyxi.gov.gr) and (2) the Greek Transparency Portal (Diavgeia, diavgeia.gov.gr).

In particular, the NSRF portal presents project-based information such as the title, budget, completion rate, related subprojects and the involved public and private organizations of a specific project. Contrarily, the Diavgeia dataset has been designed to be organization-based in the sense that every public organization has to publish all its administrative decisions. Thus, data are organized in decision types fitted to the Greek public organizations. Despite the fact that any administrative decision related to NSRF projects must be uploaded and distributed through the Diavgeia website, practically, this cannot be validated in project base because there is not a unique identifier (e.g. project code) to interconnect NSRF data to Diavgeia decisions.

The first YDS pilot is focused on bridging this gap between NSRF and Diavgeia data by identifying -through text mining- related administrative decisions for each NSRF project.

As a next step, geographical information for each project was manually added through a customized web interface for important public projects (e.g. highways and tube stations) in order to build mobile services for receiving comments and evaluations by the users.

YourDataStories Pilot #2: Official Development Assistance. The Dutch Ministry of Foreign Affairs supports many development projects across a broad range of countries, but has announced that there is a 50% cut in Dutch Official Development Assistance (ODA) starting from 2015. The Ministry also decided that at the same time it would be keen to support new projects that provide a ROI for the Dutch economy and work with particularly innovative methods. Accordingly, Dutch as well as local development NGOs are protesting. They argue that crucial humanitarian projects will come to a premature end, thus invalidating experiences and efforts built up over long time periods. This debate has attracted the attention of a team of Dutch journalists. Aim of the use case is to increase the transparency of public spending through data-driven journalism.

Initially, in pilot 2, is selected a set of 5–8 sample countries that receive development assistance from the Netherlands. Next, we retrieve detailed information on all projects run in these countries between 2011 and 2016. In step 3, we analyse information on trade relations between the sample countries and the Netherlands between 2011 and 2016. Consequently, it is retrieved detailed information on all projects run in the sample countries by Germany, France, and Denmark and on trade relations during the same period. In step 6, we made a flexibly experiment with simple charts to analyse the information above and share some expressive charts on social media. We also generated a set of research questions based on the “five Ws”, with the help of social media users and conducted a journalistic investigation.

YourDataStories Pilot #3: Cross-Europe Financial Comparability. The third pilot focus on the comparability of financial data across EU member-states, specifically looking at Ireland and Greece. Financial data from the Greek ‘Follow Public Money’ pilot will be compared with budget and spending data from the Irish national and local government, with a particular focus on construction and road infrastructure projects. Issues reported by the public via FixMyStreet.ie will also be incorporated into the data stream. Additionally, cross EU and international comparisons in public procurement are made by modelling and analysing Tenders Electronic Daily (TED) and Australian contract data.

This intelligent use of big and publicly available economic data aims to stimulate smart services for the following target users: (a) journalists to search for new stories and additional sources, (b) civil society to act more effectively in transparency and accountability issues, (c) auditors to better evaluate effectiveness and corruption in public bodies, (d) web and mobile developers to get access in cleansed and structured economic data and (e) suppliers to search for business opportunities in public procurement.

3.2 Pilot-Driven Model Extensions

The open nature of LOE allowed its direct extension with additional concepts and constructs, in order to reflect the needs of the aforementioned use cases. The following figure summarizes the model supporting the YDS platform, showcasing the involved entities and their relationships in the broader YDS context (Fig. 3).

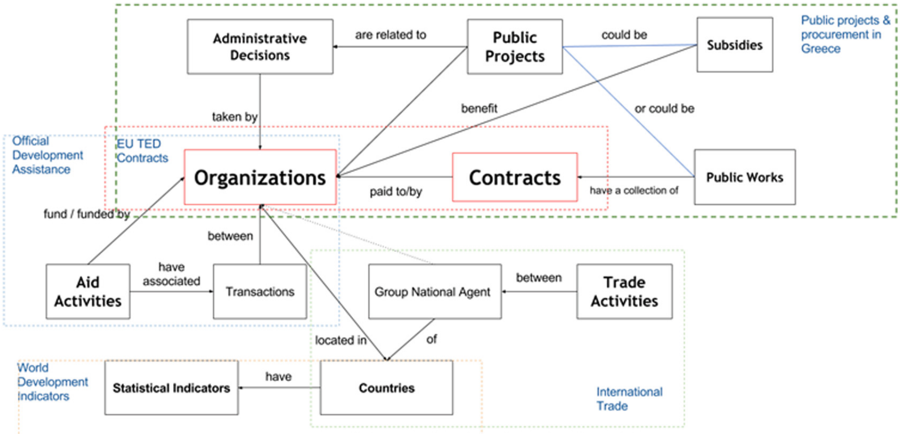

Fig. 3. YourDataStories project extension of the LOE Model

3.3 Datasets

The relevant datasets for each pilot are retrieved via dedicated, source-tailored harvesters (https://github.com/YourDataStories/harvesters) that consume the APIs or repositories of each source. The retrieved data are semi-automatically checked for correctness and consistency and are transformed into the respective RDF representations that are subsequently stored in the YourDataStories knowledge base. The following Tables 2, 3, 4 and 5 summarize the dataset size retrieved and incorporated in the YDS ecosystem.

Table 2. Diavgeia harvester results (November 2010 to November 2015)

Description	Count
Buyers	18,706
Sellers	406,955
Financial decisions	6,269,062
Produced RDF triples	166,435,634

Table 3. NSRF harvester results

Description	Count
Buyers	3,766
Sellers	17,484
Beneficiaries	126,773
Public works	12,411

Table 4. Triple count for pilot 2 data sources (graphs)

Graph	Triple count
http://yourdatastories.eu/WDI	3,816,489
http://yourdatastories.eu/OpenCorporates	10,103
http://yourdatastories.eu/WorldFactbook	8,606
http://yourdatastories.eu/taxonomies	180,695
http://yourdatastories.eu/trade/ZW	2,112,174
http://yourdatastories.eu/trade/NL	17,816,427
http://yourdatastories.eu/ODA/ZW	139,700
http://yourdatastories.eu/ODA/NL	668,919
http://yourdatastories.eu/countries	8657
http://yourdatastories.eu/regions	236

Table 5. Highest amount of EU Tender amount per country

Country	Notices	Award notices	Buyers	Sellers
FR	431,770	1,199,603	218,123	520,690
PL	176,607	997,541	26,394	265,935
UK	114,299	290,423	57,337	179,406
DE	213,520	290,956	63,286	157,393
IT	91,478	180,133	47,073	94,146
ES	93,322	198,991	37,164	91,059
BG	18,854	68,399	10,899	42,790
RO	38,418	160,304	17,213	41,814
EU	1,640,996	4,287,115	642,359	2,037,866

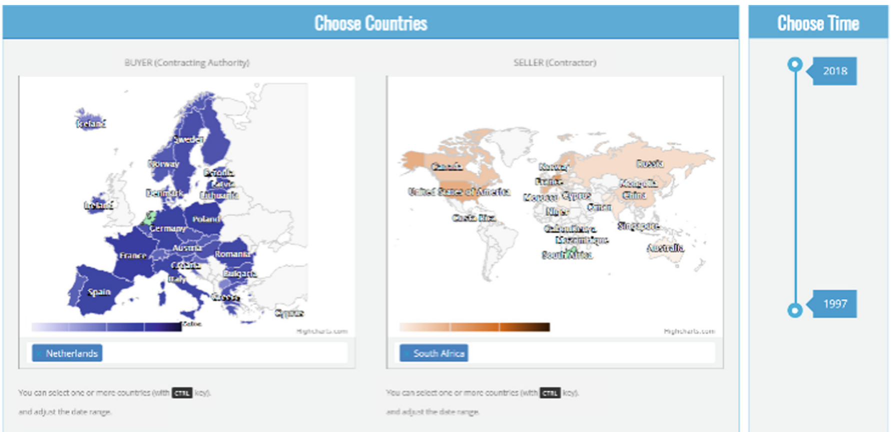

Fig. 4. YourDataStories platform interface

```
PREFIX elod: <http://linkedeconomy.org/ontology#>
PREFIX dcterms: <http://purl.org/dc/terms/>
PREFIX skos: <http://www.w3.org/2004/02/skos/core#>
PREFIX gr: <http://purl.org/goodrelations/v1#>
PREFIX foaf: < http://xmlns.com/foaf/0.1/>
PREFIX pc: < http://purl.org/procurement/public-contracts#>

select distinct ?title ?buyerName (xsd:decimal(?amount) as
?amount) ?graph
where {
{
graph ?graph {
?decision elod:hasExpenditureLine ?expLine ;
dcterms:subject ?title ; dcterms:publisher ?unit .
?buyer org:hasUnit ?unit ; skos:prefLabel ?buyerName .
?expLine elod:seller
<http://linkedeconomy.org/resource/Organization/999995379> ;
elod:amount ?ups . ?ups gr:hasCurrencyValue ?amount .
}
}
union {
graph ?graph {
?project elod:hasRelatedContract ?subproject ; elod:buyer
?buyer ; dcterms:title ?title .
?subproject elod:price ?budgetUps ; dcterms:title ?titleSub-
project ; elod:seller
<http://linkedeconomy.org/resource/Organization/999995379>.
?buyer foaf:name ?buyerName .
?budgetUps gr:hasCurrencyValue ?amount .
filter langMatches(lang(?title),'el')
}
}
union {
graph ?graph {
?awardNotice elod:documentType ""Contract award no-
tice""^^<http://www.w3.org/2001/XMLSchema#string> ;
elod:buyer ?buyer ;
elod:seller
<http://linkedeconomy.org/resource/Organization/999995379> ;
elod:contractId ?contractId .
OPTIONAL{?buyer skos:prefLabel ?buyerName .}
OPTIONAL {
?awardNotice dcterms:title ?title
} .
OPTIONAL {
?awardNotice pc:agreedPrice ?ups .
?ups gr:hasCurrencyValue ?amount
}.
}
}
}
```

Fig. 5. Combining seller information from different data collections

3.4 YourDataStories Platform and Queries

The overall LOE model, as well as, the partitions referring to the different pilot-specific domains, are openly accessible via the YDS GitHub repository (github.com/YourDataStories/ontology). The data pertaining to YDS pilots is exposed via the YDS platform (platform.yourdatastories.eu), from where the relevant stakeholders can access the underlying repositories via a SPARQL endpoint. The platform allows the immediate observation of various metrics related to the scenarios of the pilots, while the endpoint allows the direct querying of the YDS datasets (Fig. 4).

The YDS ontology distribution provides indicative queries, aiming to showcase the completeness and the added value of the model and its population. Some exemplary interesting queries are depicted in the following figures. Figure 5 depicts the query for combining seller information from different collections (graphs in the YDS repository).

Similarly, Fig. 6 presents the query for retrieving the aggregate amounts per country made available for a given CPV code.

4 Accessing and Using the LOE Model and YDS

Complete information on the LOE model and its applications can be found through the YourDataStories project website[1] and the project's GitHub repository dedicated to the development and maintenance of the ontology assets used throughout the YDS applications and use cases[2].

YDS also offers a range of dashboards for the covered datasets, where the interested users can obtain customized views for the data, essentially via a visual query editor. Direct access to the SPARQL endpoint is provided at: http://143.233.226.60:8890/, where the end user can execute the exemplary queries provided through this paper or the YDS GitHub repository, or formulate and run custom queries.

5 Conclusions

Publicly available open data are growing rapidly in quantity, but their quality can be further improved to unveil their strong potential. One particular value aspect of open data is related to their ability to address unanswered questions and provide more effective solutions in crucial policy and management issues. The LOE model is proposed as a high-level conceptualization that incorporates major economic open data, by including them in specifications that adheres to the generic CFM model. LOE, as evidenced by its support for different use cases in the context of the YDS project, is designed as the foundation for a compact but extensible common ground for journalists, professionals and public authorities to import, consume and customise open economic data.

[1] https://yourdatastories.eu/.

[2] https://github.com/YourDataStories/ontology.

```
PREFIX elod: <http://linkedeconomy.org/ontology#>
PREFIX skos: <http://www.w3.org/2004/02/skos/core#>
PREFIX gr: <http://purl.org/goodrelations/v1#>
PREFIX pc: < http://purl.org/procurement/public-contracts#>

SELECT ?countryName (sum(xsd:decimal(?agreedPriceValue)) as
?sumOfAmounts)

FROM <http://yourdatastories.eu/TEDUpdate>
FROM <http://yourdatastories.eu/TEDGreece>
FROM <http://yourdatastories.eu/TEDIreland>
FROM <http://yourdatastories.eu/taxonomies>
FROM <http://yourdatastories.eu/countries>
WHERE {
?awardNotice elod:documentType ""Contract award no-
tice""^^<http://www.w3.org/2001/XMLSchema#string> ;
elod:buyer ?buyer ;
pc:mainObject <http://linkedeconomy.org/resource/CPV/45233100-
0> ;
pc:agreedPrice ?ups .
?buyer elod:countryIsoCode ?countryIsoCode .
?ups gr:hasCurrencyValue ?agreedPriceValue .
?countryIsoCode skos:prefLabel ?countryName
}
GROUP BY ?countryName
```

Fig. 6. Cross-country CPV code budgets

6 Future Work

As stated in Sect. 2, the central ambition of the LOE model is to incorporate under a theory-wise sound conceptualization real-world data of different forms and functions. Consequently, and as financial and economic data are increasingly becoming openly available, our efforts will focus on analysing further open datasets and incorporating them in the LOE knowledge base, extending or modifying accordingly the overall model.

Another major step towards the expansion of the model and its ambition to reflect the realistic nature of public spending, is the introduction of social data in the model. Web 2.0 increasingly becomes a major source of information for facilitating the transition to the Web 3.0 vision to which LOE adheres. The initial step towards the linking of open economic data with social media will revolve around the first YDS pilot, and focus on the discovery and processing of user activity relevant to the public works incorporated in the YDS repository. To this end, the LOE model will be extended to accommodate (a) social media content and (b) its links to specific public works and their properties. The analysis of the extracted content will be modelled quantitatively (as a progress/evaluation score) and qualitatively (via the extraction of important features for the given project and the assignment of comments to the different features). A strong basis of this LOE extension is the SIOC ontology, a W3C specification and ongoing project trying to model the activities and properties of online communities, further extending the linking of the LOE schema with external conceptualizations.

Acknowledgements. Research and outcomes reported in the present manuscript have received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 645886 (YourDataStories – YDS).

References

1. Lazer, D., Pentland, A.(S.), Adamic, L., Aral, S., Barabasis, A.L., Brewer, D., Christakis, N., Contractor, N., Fowler, J., Gutmann, M., Jebara, T., King, G., Macy, M., Roy, D., Van Alstyne, M.: Life in the network: the coming age of computational social science. Science (N.Y. NY) **323**(5915), 721–723 (2009)
2. Kitchin, R.: The Data Revolution: Big Data, Open Data, Data Infrastructures and Their Consequences. Sage Publications, Thousand Oaks (2014)
3. Wolfers, J.: Big data needs economics and vice versa. Blanchard-fest Celebration. MIT, 4 June 2016
4. Smith, N.: Economics Struggles to Cope With Reality (2016)
5. Zeleti, F.A., Adegboyega, O., Curry, E.: Exploring the economic value of open government data. Gov. Inf. Q. **33**(3), 535–551 (2016)
6. World Bank: Open Data for Economic Growth. World Bank, Washington, D.C. (2014)
7. Hsinchun, C., Chiang, R.H., Storey, V.C.: Business intelligence and analytics: from big data to big impact. MIS Q. **36**(4), 1165–1188 (2012)
8. Agarwal, R., Dhar, V.: Editorial—big data, data science, and analytics: the opportunity and challenge for IS research. Inf. Syst. **25**(3), 443–448 (2014)
9. Fang, B., Zhang, P.: Big data in finance. In: Yu, S., Guo, S. (eds.) Big Data Concepts, Theories, and Applications, pp. 391–412. Springer, Cham (2016). doi:10.1007/978-3-319-27763-9_11
10. Manolea, B., Cretu, V.: Topic Report No. 2013/10: The influence of the Open Government Partnership (OGP) on the Open Data discussions. Misc, European Public Sector Information Platform (2013)
11. Tygel, A.F., Attard, J., Orlandi, F., Campos, M.L.M., Auer, S.: How much? Is not enough: an analysis of open budget initiatives. In: Proceedings of the 9th International Conference on Theory and Practice of Electronic Governance (2016)
12. Bolgov, R., Chugunov, A., Filatova, O.: Open budgets in Eurasian countries: Cloacae or Klondike? In: Conference for E-Democracy and Open Government (2015)
13. Brusa, G., Caliusco, M., Chiotti, O.: A process for building a domain ontology: an experience in developing a government budgetary ontology. In: Proceedings of the Second Australasian workshop on Advances in Ontologies (2006)
14. de Oliveira Araújo, L.S., Santos, M.T., Silva, D.A.: The Brazilian federal budget ontology: a semantic web case of public open data. In: Proceedings of the 7th International Conference on Management of computational and collective intElligence in Digital EcoSystems (2015)
15. Höffner, K., Martin, M., Lehmann, J.: LinkedSpending: OpenSpending becomes Linked Open Data. Semant. Web **7**(1), 95–104 (2015)
16. de Santana, M., da Silva Craveiro, G.: Challenges and requirements for the standardisation of open budgetary data in the Brazilian public administration. In: GI-Jahrestagung (2013)
17. Vafopoulos, M., Meimaris, M.: Weaving the economic Linked Open Data. In: 7th International Workshop on Semantic and Social Media Adaptation and Personalization (SMAP) (2012)
18. Alvarez-Rodríguez, J.M., Labra-Gayo, J.E., de Pablos, P.O.: New trends on e-Procurement applying semantic technologies: current status and future challenges. Comput. Ind. **65**(5), 800–820 (2014)

19. Open Contracting Partnership: Open Contracting Data Standard (OCDS), version 1.0 (2015)
20. Vafopoulos, M., Meimaris, M., Papantoniou, A., Alexiou, G., Avraam, I., Xidias, I., Loumos, V.: Publicspending.gr : Greek public spending in a nutshell. In: Samos Summit on Open Data and Interoperability (2012)
21. Vafopoulos, M., Rodríguez, J., Meimaris, M., Xidias, I., Klonaras, M., Vafeiadis, G.: Insights in global public spending. In: 13th International Conference on Knowledge Management and Knowledge Technologies and 9th International Conference on Semantic Systems, Messe Congress Graz, Austria, 2013
22. Vafopoulos, M., Meimaris, M., Anagnostopoulos, I., Papantoniou, A., Xidias, G.A.I., Loumos, V.: Public spending as LOD: the case of Greece. Semant. Web J., 155–164 (2014)
23. Taggart, Ch., Peltola, V.: OpenCorporates: why we're crying out for this database of companies. The Guardian (2011). http://bit.ly/2ilSwdD. Accessed 10 Oct 2017

Unsupervised Keyword Extraction Using the GoW Model and Centrality Scores

Elissavet Batziou[1], Ilias Gialampoukidis[2(✉)], Stefanos Vrochidis[2], Ioannis Antoniou[1], and Ioannis Kompatsiaris[2]

[1] School of Mathematics, Aristotle University of Thessaloniki, Thessaloniki, Greece
[2] Centre for Research and Technology Hellas, Information Technologies Institute, Thessaloniki, Greece
heliasgj@iti.gr

Abstract. Nowadays, a large amount of text documents are produced on a daily basis, so we need efficient and effective access to their content. News articles, blogs and technical reports are often lengthy, so the reader needs a quick overview of the underlying content. To that end we present graph-based models for keyword extraction, in order to compare the Bag of Words model with the Graph of Words model in the keyword extraction problem. We compare their performance in two publicly available datasets using the evaluation measures Precision@10, mean Average Precision and Jaccard coefficient. The methods we have selected for comparison are grouped into two main categories. On the one hand, centrality measures on the formulated Graph-of-Words (GoW) are able to rank all words in a document from the most central to the less central, according to their score in the GoW representation. On the other hand, community detection algorithms on the GoW provide the largest community that contains the key nodes (words) in the GoW. We selected these methods as the most prominent methods to identify central nodes in a GoW model. We conclude that term-frequency scores (BoW model) are useful only in the case of less structured text, while in more structured text documents, the order of words plays a key role and graph-based models are superior to the term-frequency scores per document.

Keywords: Keyword-based search · Topic-based filtering · Graph-based models · Graph of words · Centrality measures · Community detection

1 Introduction

Textual information is all around us (smartphones, WWW, social media, etc.), involving large streams of information with content that needs to be accessed quickly. At the level of a single document, reading lengthy text documents is a time consuming process that needs to be assisted by a keyword extraction mechanism, in order to provide the reader a quick overview of the main topics

I. Kompatsiaris et al. (Eds.): INSCI 2017, LNCS 10673, pp. 344–351, 2017.
https://doi.org/10.1007/978-3-319-70284-1_26

of the text document. Keyword extraction needs to be an automatic process, assisted by efficient and effective text representations that exploit graph models.

The methods that have been used for keyword extraction on the graph of words are grouped into two main categories. Firstly, centrality measures and more general centrality-based scores (transitivity, coreness) are employed, being able to rank all words in a document from the most central to the less central, according to their score in the GoW representation. Secondly, community detection algorithms on the GoW provide the largest community that contains the key nodes (words) in the GoW. Graph-based keyword extraction methods are reported in [1].

Betweenness centrality has been used in the context of keyword extraction [1], as well as the closeness centrality [2], the degree centrality [3], Eigenvector centrality [4] and PageRank [5]. In addition, eccentricity [6] and coreness, transitivity (known also as clustering coefficient) and Term-Frequency (TF) scores have been examined in keyword extraction [3].

The largest community of the graph of words may also be extracted to provide a group of words as the most representative ones in the text document. This approach have been discussed in [7], where the extraction of the key-community of words is done using the edge betweenness modularity maximization method.

The purpose of this paper is to review unsupervised graph-based models and to compare them in two public annotated collections. We propose and examine alternative centrality-based methods to extract keywords from the Graph of Words (GoW) model, which is an extension of the Bag of Words (BoW) representation model.

Our paper is structured as follows. In Sect. 2 we present the BoW and GoW text representation models and in Sect. 3 we additionally provide centrality measures and community detection approaches for the extraction of keywords from text documents, when they are combined with a graph of words. In Sect. 4 we examine which method performs better in public datasets and finally in Sect. 5 we conclude our paper.

2 BoW and GoW Models

We describe and apply the GoW model in the keyword extraction problem and we compare its performance with the BoW model, as obtained from the most frequent terms in a document.

2.1 BoW Model

The Bag-of-words (BoW) model is a text representation which have been used in Natural Language Processing (NLP) and in Information Retrieval (IR). In this model, text is represented as a bag which contains all text's words, free from grammar and word order. Word's multiplicity is the number of occurrences of a word in a document, known also as term frequency (tf):

$$\mathrm{tf}_{wd} = \frac{n_{wd}}{n_d} \tag{1}$$

where n_{wd} is the number of occurrences of word w in document d and n_d is the number of words in document d.

Term frequency (tf) scores are weighted by the inverse document frequency, to put less weight in words that appear in many documents. The tfidf scores are defined as:

$$\text{tf-idf}_{wd} = \frac{n_{wd}}{n_d} \log \frac{N}{n_w} \tag{2}$$

where N is the total number of documents in the database and n_w is the number of occurrences of word w in the whole database.

2.2 GoW Model

Graph of words (GoW) is the representation of a text document as an unweighted graph [8], where its nodes represent terms (words). Given a window of N successive words in a document, all terms in the window are mutually linked and each edge represents the co-occurrence of a pair of terms in the window set. Links on the graph representation of a text document are provided by bi-grams and/or tri-grams, according to the size of the considered window of $N = 2$ or $N = 3$ respectively. Contrary to the BoW representation, the GoW model exploits n-grams to formulate the graph of words and, moreover, keeps the complex structure of the interdependencies among all n-grams. Using for example only the word frequency in a text (unigrams), the model will not reveal the fact that after a name follows a verb in the text, but the n-grams keep this information, as shown in Fig. 1.

(a) $N = 2$ (b) $N = 3$

Fig. 1. Graph of words for $N = 2$ and $N = 3$ on the text "The international conference on internet science aims at progressing and investigating on topics of high relevance with internet's impact on society, governance, and innovation. It focuses on the contribution and role of internet science on the current..."

3 Keyword Extraction Using the GoW Model and Centrality Measures

We examine the performance of the following centrality measures in the keyword extraction problem, as recent centrality measures that have been introduced in Statistical Mechanics [9] or Security Informatics [10], namely Mapping Entropy and Mapping Entropy Betweenness (MEB), respectively.

Let G be the graph of words, where $\mathcal{N}(n_k)$ denotes the neighborhood of the node n_k. We also propose a novel centrality measure, motivated by Mapping Entropy [9] and MEB [10], as follows:

$$\mathrm{MEC}_k = -CC_k \sum_{n_i \in \mathcal{N}(n_k)} \log CC_i \tag{3}$$

where CC_i is the closeness centrality of node n_i. Hence, the proposed centrality measure is called Mapping Entropy Closeness (MEC).

The community detection approach for keyword extraction [7], is based on the maximization of modularity. In the following experiments (Sect. 4) we moreover examine the performance of the largest detected community of words, in the GoW representation, as extracted by one of the following approaches:

- Fast greedy (modularity maximization) [11]
- Infomap (codelength minimization) [12]
- Label Propagation [13]
- Louvain (modularity maximization) [14]
- Walktrap (random walks) [15]

The results are presented in the following section.

4 Experimental Comparison

In this section we examine which mode is more suitable to the keyword extraction problem, in two diverse public datasets. We involve all methods that have been discussed in Sect. 2 in our experimental comparison, which is done under the evaluation measures Precision at 10 (P@10) and Average Precision, which are popular in IR tasks. Moreover, the Jaccard index is able to measure the similarity between the ground-truth list of keywords and the keywords that are extracted by each method.

4.1 Dataset Description

The datasets we have selected for comparison are, firstly, the Fao780[1] dataset which contains 779 documents and the CiteULike180[2] dataset with 183 text

[1] https://github.com/zelandiya/keyword-extraction-datasets.
[2] https://github.com/snkim/AutomaticKeyphraseExtraction.

documents, tagged by 152 taggers. The CiteULike dataset has 183 publications crawled from CiteULike, and keywords assigned by different CiteULike users who saved these publications. The other dataset, FAO780, has 779 FAO publications with Agrovoc terms from official documents of the Food and Agriculture Organization of the United Nations (FAO).

4.2 Settings

Firstly, we remove punctuation and we transform all letters to lowercase. Numbers are also removed, as well as the English stopwords, which are common words that are repeated (e.g. "the", "a", "and") without adding meaning to the document, known as the SMART[3] stopwords list. Moreover, we stem each word, i.e. we remove the ending of the word, so as to keep only the word's stem. Afterwards, we construct the graph of words, which has as nodes the words of our document. Two nodes take link if a word follows the other, i.e. any two terms of a bi-gram ($N = 2$) are connected. We also examine the performance of the keyword extraction problem, by linking the terms of tri-grams ($N = 3$).

In all datasets, we keep the top-20 keywords for each selected centrality score (Betweenness, Closeness, Degree, Eigenvector, Page Rank, Mapping Entropy, MEB, MEC, Coreness, Transitivity, Eccentricity) and for the top-20 most frequent terms (TF scores). In the case of the most informative community of the constructed graph of words, we use five prominent community detection algorithms (Fast greedy, Infomap, Label Prop, Louvain and Walktrap).

4.3 Results

The GoW model is superior to the BoW representation, in the case of structured text, as shown in Table 1 and in Table 2 for a window of size $N=2$ and $N = 3$, respectively. FAO documents have more unstructured text than CiteULike documents, where we present two sample text documents from these datasets in Fig. 2. In the case of structured text (CiteULike), we observe that the GoW representation performs better than the simple statistical term frequency scores. On the other hand, in the FAO dataset, term frequency scores count the most frequent words and are able to identify the most critical words in each document. In structured text, the order of words is very important because links are added between a word and its N successive words.

Given the GoW representation, we observe that when $N = 3$ the results are better than the case of $N = 2$, where N is the number of successive words that are linked. However, the linking of more words than $N = 3$ successive words, makes the graph of words almost complete, so centralities become identical and the graph has only one community (all the graph).

Among the centrality measures, closeness centrality performs better than the other measures. In the case of $N = 2$, Mapping Entropy Betweenness centrality

[3] http://jmlr.csail.mit.edu/papers/volume5/lewis04a/a11-smart-stop-list/english.stop.

Table 1. Jaccard, Average Precision and P@10 results for linking $N = 2$ successive words.

N = 2	CiteULike180			FAO780		
Method	Jaccard	Av Prec	P@10	Jaccard	Av Prec	P@10
Betweenness	0.1531 ± 0.0598	0.3795 ± 0.1404	0.3486 ± 0.1398	0.1619 ± 0.0734	0.3459 ± 0.1500	0.3112 ± 0.1473
Closeness	0.1531 ± 0.0622	0.3890 ± 0.1425	0.3552 ± 0.1413	0.1656 ± 0.0781	0.3565 ± 0.1547	0.3212 ± 0.1540
Degree	0.1566 ± 0.0611	0.3842 ± 0.1390	0.3492 ± 0.1410	0.1671 ± 0.0777	0.3533 ± 0.1538	0.3208 ± 0.1508
Eigenvector	0.1446 ± 0.0659	0.3606 ± 0.1453	0.3525 ± 0.1421	0.1649 ± 0.0792	0.3526 ± 0.1570	0.3158 ± 0.1549
Page Rank	0.0508 ± 0.0313	0.3831 ± 0.1399	0.3492 ± 0.1410	0.1669 ± 0.0772	0.3488 ± 0.1530	0.3173 ± 0.1503
Mapping Ent	0.1557 ± 0.0613	0.3821 ± 0.1394	0.3519 ± 0.1406	0.1669 ± 0.0780	0.3515 ± 0.1533	0.3191 ± 0.1502
MEB	0.1598 ± 0.0625	0.3860 ± 0.1378	0.3530 ± 0.1354	0.0674 ± 0.0451	0.1762 ± 0.1180	0.1469 ± 0.1009
MEC	0.1567 ± 0.0622	0.3839 ± 0.1389	0.3503 ± 0.1402	0.0678 ± 0.0460	0.1753 ± 0.1178	0.1477 ± 0.1009
Coreness	0.1098 ± 0.5110	0.2857 ± 0.1364	0.3508 ± 0.1568	0.0839 ± 0.0487	0.1802 ± 0.0994	0.2855 ± 0.1556
Transitivity	0.0000 ± 0.0000	0.0182 ± 0.0469	0.0164 ± 0.0426	0.0067 ± 0.0154	0.0221 ± 0.0559	0.0171 ± 0.0422
Eccentricity	0.0015 ± 0.0062	0.0026 ± 0.0157	0.0027 ± 0.0163	0.0003 ± 0.0033	0.0004 ± 0.0054	0.0004 ± 0.0062
TF score	0.1613 ± 0.0648	0.3877 ± 0.1421	0.3530 ± 0.1386	0.1781 ± 0.0843	0.3725 ± 0.1603	0.3392 ± 0.1614
Fast greedy	0.0215 ± 0.0164	0.0649 ± 0.0500	0.1656 ± 0.1459	0.0100 ± 0.0116	0.0297 ± 0.0303	0.1163 ± 0.1114
Infomap	0.0402 ± 0.0248	0.1258 ± 0.0762	0.2749 ± 0.1770	0.0205 ± 0.0220	0.0586 ± 0.0581	0.2258 ± 0.1462
Label Prop	0.0158 ± 0.0088	0.0411 ± 0.0203	0.2754 ± 0.1693	0.0074 ± 0.0069	0.0219 ± 0.0153	0.2100 ± 0.1420
Louvain	0.0193 ± 0.0167	0.0600 ± 0.0538	0.1421 ± 0.1415	0.0107 ± 0.0130	0.0320 ± 0.0359	0.0992 ± 0.1054
Walktrap	0.0332 ± 0.0171	0.0941 ± 0.0459	0.3060 ± 0.1846	0.0176 ± 0.0173	0.0504 ± 0.0412	0.2144 ± 0.1439

Table 2. Jaccard, Average Precision and P@10 results for linking $N = 3$ successive words.

N = 3	CiteULike180			FAO780		
Method	Jaccard	Av Prec	P@10	Jaccard	Av Prec	P@10
Betweenness	0.1609 ± 0.0633	0.3854 ± 0.1431	0.3519 ± 0.1441	0.1671 ± 0.0748	0.3568 ± 0.1505	0.3213 ± 0.1504
Closeness	0.1658 ± 0.0617	0.4034 ± 0.1447	0.3776 ± 0.1490	0.1731 ± 0.0819	0.3678 ± 0.1560	0.3326 ± 0.1558
Degree	0.1648 ± 0.0621	0.3993 ± 0.1406	0.3661 ± 0.1404	0.1744 ± 0.0806	0.3671 ± 0.1543	0.3304 ± 0.1532
Eigenvector	0.1542 ± 0.0629	0.3791 ± 0.1445	0.3448 ± 0.1428	0.1711 ± 0.0818	0.3662 ± 01589	0.3291 ± 0.1590
Page Rank	0.1645 ± 0.0662	0.3982 ± 0.1401	0.3678 ± 0.1395	0.1740 ± 0.0807	0.3641 ± 0.1542	0.3286 ± 0.1530
Mapping Ent	0.1644 ± 0.0632	0.3974 ± 0.1404	0.3650 ± 0.1394	0.1746 ± 0.0807	0.3662 ± 0.1544	0.3295 ± 0.1540
MEB	0.1638 ± 0.0619	0.3963 ± 0.1397	0.3661 ± 0.1435	0.1723 ± 0.0776	0.3627 ± 0.1527	0.3293 ± 0.1530
MEC	0.1648 ± 0.0636	0.3886 ± 0.1407	0.3683 ± 0.1402	0.1745 ± 0.0803	0.3671 ± 0.1544	0.3295 ± 0.1527
Coreness	0.1066 ± 0.0481	0.2637 ± 0.1208	0.3694 ± 0.1682	0.075 ± 0.0440	0.1595 ± 0.0848	0.2796 ± 0.1542
Transitivity	0.0015 ± 0.0062	0.0025 ± 0.0161	0.0022 ± 0.0147	0.0001 ± 0.0050	0.0015 ± 0.0130	0.0014 ± 0.0118
Eccentricity	0.0016 ± 0.0067	0.0022 ± 0.0124	0.0033 ± 0.0179	0.0006 ± 0.0045	0.0010 ± 0.0090	0.0006 ± 0.0080
TF score	0.1613 ± 0.0648	0.2637 ± 0.1208	0.3530 ± 0.1386	0.1781 ± 0.0843	0.3725 ± 0.1603	0.3392 ± 0.1614
Fast greedy	0.0196 ± 0.0146	0.0565 ± 0.0399	0.1792 ± 0.1475	0.0086 ± 0.0098	0.0255 ± 0.0257	0.1167 ± 0.1169
Infomap	0.0283 ± 0.0167	0.0865 ± 0.0490	0.2995 ± 0.1903	0.014 ± 0.0145	0.0407 ± 0.0393	0.2248 ± 0.1423
Label Prop	0.0151 ± 0.0077	0.0394 ± 0.0181	0.2689 ± 0.1696	0.0072 ± 0.0066	0.0216 ± 0.0147	0.2089 ± 0.1412
Louvain	0.0160 ± 0.0154	0.0464 ± 0.0444	0.1235 ± 0.1294	0.0098 ± 0.0111	0.0288 ± 0.0298	0.1141 ± 0.1166
Walktrap	0.0280 ± 0.0166	0.0809 ± 0.0436	0.2891 ± 0.1895	0.0140 ± 0.0136	0.0414 ± 0.0347	0.1979 ± 0.1418

has larger Jaccard index than all other methods. Among the community detection approaches, the Infomap communities contain the most important words on average and therefore obtain higher Jaccard, Average Precision and P@10.

Community detection approaches are not superior to centrality scores, in all cases examined. Our proposed Mapping Entropy Closeness (MEC) centrality measure is the second most performing keyword extraction approach, in the case of Jaccard index, following the Mapping Entropy Betweenness (MEB) scores.

FAO sample text

Where to purchase FAO publications locally - Points de vente des publications de la FAO - Puntos de venta de publicaciones de la FAO
· ANGOLA
Empresa Nacional do Disco e de Publicaηues, ENDIPU-U.E.E.
Rua Cirilo da Conceiηγo Silva, N° 7
C.P. N° 1314-C, Luanda
· ARGENTINA
Librerva Agropecuaria
Pasteur 743, 1028 Buenos Aires
Oficina del Libro Internacional
Av. Cαrdoba 1877, 1120 Buenos Aires
E-mail: olilibro@satlink.com
· AUSTRALIA
Hunter Publications
P.O. Box 404, Abbotsford, Vic. 3067
Tel.:(03) 9417 5361
Fax: (03) 914 7154
E-mail: jpdavies@ozemail.com.au
· AUSTRIA
Gerold Buch & Co.
Weihburggasse 26, 1010 Vienna

CiteULike sample text

The study of networks pervades all of science, from neurobiology to statistical physics. The most basic issues are structural: how does one characterize the wiring diagram of a food web or the Internet or the metabolic network of the bacterium Escherichia coli? Are there any unifying principles underlying their topology? From the perspective of nonlinear dynamics, we would also like to understand how an enormous network of interacting dynamical systems -- be they neurons, power stations or lasers -- will behave collectively, given their individual dynamics and coupling architecture. Researchers are only now beginning to unravel the structure and dynamics of complex networks. Networks are on our minds nowadays. Sometimes we fear their power -- and with good reason. On 10 August 1996, a fault in two power lines in Oregon led, through a cascading series of failures, to blackouts in 11 US states and two Canadian provinces, leaving about 7 million customers without power for up to 16 hours1. The Love Bug worm, the worst computer attack to date, spread over the Internet on 4 May 2000 and inflicted billions of dollars of damage worldwide. In our lighter moments we play parlour games about connectivity.

Fig. 2. Sample from FAO and CiteULike text documents.

5 Conclusion

We used graph-based models to extract keywords from text documents. We examined the performance of 17 keyword extraction techniques based on centrality measures and community detection approaches on the graph of words. We observed that in the case of structured text the GoW representation performs better than the simple statistical term frequency scores. On the other hand, term frequency scores were able to identify the most critical words in each document where text is less structured. We also proposed the Mapping Entropy Closeness (MEC) centrality measure which is the second most performing keyword extraction approach, in the case of Jaccard index, following the Mapping Entropy Betweenness (MEB) scores. Centrality scores outperform community detection approaches in keyword extraction in all datasets examined.

Acknowledgements. This work was supported by the projects H2020-645012 (KRISTINA) and H2020-700024 (TENSOR), funded by the European Commission.

References

1. Beliga, S., Meštrović, A., Martinčić-Ipšić, S.: An overview of graph-based keyword extraction methods and approaches. J. Inf. Organ. Sci. **39**(1), 1–20 (2015)
2. Abilhoa, W.D., De Castro, L.N.: A keyword extraction method from twitter messages represented as graphs. Appl. Math. Comput. **240**, 308–325 (2014)
3. Lahiri, S., Choudhury, S.R., Caragea, C.: Keyword and keyphrase extraction using centrality measures on collocation networks. arXiv preprint arXiv:1401.6571 (2014)
4. Boudin, F.: A comparison of centrality measures for graph-based keyphrase extraction. In: International Joint Conference on Natural Language Processing (IJCNLP), pp. 834–838 (2013)

5. Tsatsaronis, G., Varlamis, I., Nørvåg, K.: Semanticrank: ranking keywords and sentences using semantic graphs. In: Proceedings of the 23rd International Conference on Computational Linguistics, pp. 1074–1082. Association for Computational Linguistics (2010)
6. Xie, Z.: Centrality measures in text mining: prediction of noun phrases that appear in abstracts. In: Proceedings of the ACL Student Research Workshop, pp. 103–108. Association for Computational Linguistics (2005)
7. Grineva, M., Grinev, M., Lizorkin, D.: Extracting key terms from noisy and multitheme documents. In: Proceedings of the 18th international conference on World wide web, pp. 661–670. ACM (2009)
8. Rousseau, F., Vazirgiannis, M.: Graph-of-word and TW-IDF: new approach to ad hoc IR. In: Proceedings of the 22nd ACM International Conference on Information & Knowledge Management, pp. 59–68. ACM (2013)
9. Nie, T., Guo, Z., Zhao, K., Lu, Z.M.: Using mapping entropy to identify node centrality in complex networks. Phys. A **453**, 290–297 (2016)
10. Gialampoukidis, I., Kalpakis, G., Tsikrika, T., Vrochidis, S., Kompatsiaris, I.: Key player identification in terrorism-related social media networks using centrality measures. In: European Intelligence and Security Informatics Conference (EISIC 2016), pp. 17–19. August (2016)
11. Clauset, A., Newman, M.E.J., Moore, C.: Finding community structure in very large networks. Phys. Rev. E **70**(6), 066111 (2004)
12. Rosvall, M., Bergstrom, C.T.: Maps of random walks on complex networks reveal community structure. Proc. Nat. Acad. Sci. **105**(4), 1118–1123 (2008)
13. Raghavan, U.N., Albert, R., Kumara, S.: Near linear time algorithm to detect community structures in large-scale networks. Phys. Rev. E **76**(3), 036106 (2007)
14. Blondel, V.D., Guillaume, J.L., Lambiotte, R., Lefebvre, E.: Fast unfolding of communities in large networks. J. Stat. Mech: Theory Exp. **2008**(10), P10008 (2008)
15. Pons, P., Latapy, M.: Computing communities in large networks using random walks. J. Gr. Algorithms Appl. **10**(2), 191–218 (2006)

Implicit Interaction Through Machine Learning: Challenges in Design, Accountability, and Privacy

Donald McMillan(✉)

Stockholm University, Stockholm, Sweden
donald.mcmillan@dsv.su.se

Abstract. Implicit Interaction takes advantage of the rise of predictive algorithms, trained on our behaviour over weeks, months and years, and employs them to streamline our interactions with devices from smartphones to Internet connected appliances. Implicit Interaction provides users the advantage of systems that learn from their actions, while giving them the feedback and controls necessary to both understand and influence system behaviour without having to rely on an application for every connected device. This is an active area of research and as such presents challenges for interaction design due, in part, to the use of user-facing machine learning algorithms. This paper discusses the challenges posed by designing in accountability for system actions and predictions, the privacy concerns raised by both the sensing necessary to power these predictions and in how the predictions and systems actions themselves can expose behavioural patterns, and the challenges inherent in designing for the reality of machine learning techniques rather than the hype.

Keywords: Implicit interaction · Internet of things · Machine learning · Privacy · Interaction design

1 Introduction

The Internet of Things (IoT) presents several visions of the future. One, which extrapolates from current practice of interacting with nascent IoT technology, is of providing apps for everything [21] to be installed, updated, and learned alongside the tasks of installing, updating, and learning to use the IoT devices themselves. A competing vision is that of Implicit Interaction.

Implicit interactions stay in the background, thriving on analysis of speech, movement, and other contextual data, avoiding unnecessarily disturbing us or grabbing our attention. When we turn to them, depending on context and functionality, they either shift into an explicit interaction-engaging us in a classical interaction dialogue informed by the analysis of the context at hand, alternatively they continue to engage us implicitly using entirely different modalities that do not require an explicit dialogue that is through the way we move or engage in other tasks, the smart objects respond to us. For example, one form of implicit interaction is when mobile phones listen to surrounding conversation and continuously adapt to what might be a relevant starting

I. Kompatsiaris et al. (Eds.): INSCI 2017, LNCS 10673, pp. 352–358, 2017.
https://doi.org/10.1007/978-3-319-70284-1_27

point once the user decides to turn to it [17]. As the user activates the mobile, we can imagine how the search app already has search terms from the conversation inserted, the map app shows places discussed in the conversation, or if the weather was mentioned and the person with the mobile was located in their garden, the gardening app may have integrated the weather information with the sensor data from the humidity sensor in your garden to provide a relevant starting point. This is of course only possible through providing massive data sets and making continuous adaptations to what people say, their indoor and outdoor location, their movements and any smart objects in that environment – thriving off the whole ecology of artefacts, people and their practices. The implicit interaction paradigm, however, presents unique challenges when dealing with accountability, privacy, and control. This paper discusses each of these in turn, and sets up a series of challenges for researchers in this field.

2 Background

Implicit Interaction builds on the history of intelligent agents [20], behavioural inference [12] and motion sensing [2]. Implicit interactions stay in the background [6, 13, 18], thriving on data analysis of speech [17] and movements [16]. On a trajectory towards this vision, HCI researchers have used machine learning in interaction in a variety of, from turning the body into an interactive interface [10] to creating adaptive interfaces that automate and facilitate user's tasks reduce [15]. Machine learning has been used to understand the routines of users [5] including their interruptability during certain tasks [8, 9], in order to better support interaction. Better understanding the user themselves, rather than their routines, has been used to provide systems that can detect depression [7], recommend products and services [11], and produce intelligent tutors [1].

However designing with machine learning is an ongoing challenge and area of research. Dove *et al.* surveyed UX practitioners [16] noting problems in understanding the capabilities of Machine Learning, and integrating it into design practices such as prototyping. Systems based on machine learning may be unable to show understanding of users' intent, leading them to be perceived as useless and unintuitive [22], and in studies of self-driving vehicles human behaviour has been shown to be responsive to real and perceived levels of control [4].

3 Privacy and Accountability

The vision of Implicit Interaction is built on collection and training on large amounts of personal data, and making the sensing and collection of this data happen in a privacy sensitive manner is a challenge facing any number of fields concerned with human subjects looking to harness the growing power of machine learning. The recent court case in the USA surrounding data collected, possibly inadvertently, by the Amazon Echo home assistant [19] highlights one issue with collecting large amounts of data not explicitly directed at controlling the devices provided. Building on Privacy by Design recommendations [14], we can suggest that systems should not store raw sensor data, such as the unrecognized audio recordings under scrutiny in the Echo case above, but this in turn causes problems for the future of said systems. The cases where the system was unable to understand the users' intention are the cases where the raw data can provide valuable input to improve the algorithm, not training the system using these pieces of data would require significantly more effort in collecting and maintaining a training data set. Even without the raw sensor data, information about the actions of users could be retrieved by interrogating the machine learning algorithm. By comparing the output of the underlying algorithm taken in its base state and the trained state for a particular user when presented with different inputs it would be possible to provide a probability that a certain input pattern had been presented to that algorithm. While this could not provide someone with new data it holds the possibility that such techniques would allow law enforcement, or other actors, to determine with some certainty if a certain action had been taken or phrase had been spoken in a location.

More than just the data, the actions of the system can cause problems with privacy. The actions are a reflection of the dominant behavior of the user or users, and in a system where the goal is to pre-empt users to adjust the environment to them without any explicit commands this could expose behavior that the user would rather keep private. The precursor to this problem has already been seen, when Target (a large retailer in the USA) through their loyalty card system predicted a teen's pregnancy and, with targeted adverts and discounts, inadvertently informed the new grandfather before his daughter had informed him that she was pregnant[1]. This calls for research into how best to give control of the learning function of the algorithms that will watch our everyday lives back to the users being watched, and it is two fold. On one hand, users would benefit from the ability to pause the learning of systems in their home in unusual situations (such as renting their home to strangers on AirBnB) or in situations they would not want reflected in the presentation of self that the automatic actions of a system trained on their actions would become. On the other hand, research is needed in ways to allow users to interrogate algorithms and make the current actions of the system accountable to past actions of the users. Not only would this allow 'explaining away' embarrassing situations caused by such systems, but it would provide understanding of the effect their actions have on the learning algorithm necessary to consciously change their behavior to result in the system behaving in a manner they choose.

As the learning algorithms watch user activity over time, finding and reinforcing predictive correlations between behaviour and the user's interaction with the actuators under the control of the system, understanding why the system behaves as is does can require a level of reflection not necessary in short-loop interactions. The challenge here is not only providing an understandable representation of the progression of training in relation to the recorded actions of the users, it also requires research to extract and understand the current state of training of the learning algorithm which in itself is a current topic of research. Indeed, the 'black box' nature of many deep learning algorithms hinders the accountability of the algorithms significantly. In the vision of Implicit Interaction there would be cascade of learning algorithms, with those closest to

[1] https://www.forbes.com/sites/kashmirhill/2012/02/16/how-target-figured-out-a-teen-girl-was-pregnant-before-her-father-did/.

the user detecting simple individual actions (such as room occupancy, wakefulness, or low level activities like reading a book), which would in turn be used as input to those dedicated to higher level understanding of activity over time, and ideally these would both be inputs into a further stage of machine learning algorithms looking at user intent and supporting higher level longer term goals for user engagement.

In such a tiered system, it would be possible to store and display the inputs that had been detected by the lower level algorithms alongside the results of the higher-level ones, providing at least some semblance of meaning and accountability to system actions. However, before research into making the underlying algorithms accountable bears fruit this will provide only a small window through which the user must make sense of the actions of the system and how they can influence it.

4 Implicit Control

Controlling a system via a machine learning algorithm will present a number of challenges in itself. In challenging the accountability problem above, the users must be informed through design of the possibilities and limitations with machine learning algorithms, and before this can be attempted designers must develop the understanding necessary to design control mechanisms that fit within these limitations. This presents a challenge as the current narrative on machine learning is described in the Gartner Hype Cycle[2] as being at the “Peak of Inflated Expectations” – meaning that many of the assumptions about machine learning disseminated outside the community of practitioners and researchers working directly with these algorithms can be taken with a pinch of salt.

The most important ‘Inflated Expectation’ sold in relation to work on Implicit Interaction is that given ‘enough’ data a system could learn to understand rare and infrequent contexts of human behaviour. Machine learning excels in categorising complex data into common and relatively balanced categories, but when the data ventures into the realm of ‘Imbalanced Domains’ [3] there are a number of extra challenges in developing algorithms that can effectively categorise the data. There are two distinct challenges in imbalanced domains that should be understood by those designing for and with machine learning, but in order to explain them the default behaviour of a machine learning algorithm over a data set must also be described.

Take, for example, a data set consisting of 1 million tweets relating to a political party. A machine learning algorithm trained to determine if the tweets are positive or negative in sentiment would be shown a subset of these tweets categorised by an expert. Each one of these training tweets is fed into the algorithm and it provides its guess as to if the tweet is positive or negative. As this is during the training of the algorithm, the outcome is checked against the expert categorisation. If the machine matches the expert then the algorithm parameters (which differ from one method to the next) are reinforced, if the machine gets it wrong then the parameters are changed in the other direction. This is done over thousands and thousands of tweets with the overall

[2] http://www.gartner.com/newsroom/id/3412017.

accuracy a goal for the system. Carrying on this overly simplistic abstraction we can see that if the data set consisted of 20% of positive tweets and 80% negative tweets then the algorithm could return an 80% accuracy simply by categorising all data as negative – something that beginner's implementations of machine learning algorithms often do.

At this level there are any number of techniques, beyond simply looking in more detail at the type of error, to ensure that this doesn't happen in production systems yet the underlying problem is still the same when the number of categories increases and the size of the small categories decreases. The two problems with Imbalanced Domains, rephrased from [3], can be described as: (1) It is more important for the user to get accurate results from some categories in the data than it is from others, and (2) the cases that are more important to the user are under-represented in the data.

In the context of home automation, it may be more important to the user for the system to recognise rare occurrences accurately (a burglary or an acute illness, for example) at the expense of occasionally misrecognising the intention to open the window as closing the curtains.

Given enough information about the preferences of the user the choice of algorithm and the learning method can be adjusted to improve accuracy, but getting that information from the user is difficult. This is a major challenge for interaction design: Providing methods and metaphors to allow users to understand and influence the learning algorithms possibly by exposing the imbalances in the data in order to provide a counterbalance of user preferences.

5 Conclusion

This paper presents challenges faced in HCI by the advances in machine learning and their ongoing incorporation into user-facing systems. The Implicit Interaction project described in brief above is one of many research initiatives striving to bridge the gaps in knowledge and practice between users, machine learning experts, interaction designers, and the machine learning algorithms themselves. By increasing the awareness of the problems presented by the data that feeds the algorithms, the social and societal implications of the inferences that they make, and the disconnect between expectations of their abilities and the realities of implementing such algorithms this paper provides a starting point for more nuanced discussion of the ever increasing influence that machine learning algorithms have in our everyday lives.

References

1. Aleven, V., Mclaren, B.M., Sewall, J., Koedinger, K.R.: A new paradigm for intelligent tutoring systems: example-tracing tutors. Int. J. Artif. Intell. Ed. **19**(2), 105–154 (2009)
2. Bachynskyi, M., Palmas, G., Oulasvirta, A., Weinkauf, T.: Informing the design of novel input methods with muscle coactivation clustering. ACM Trans. Comput. Hum. Interact. (TOCHI) **21**(6), 30 (2015)

3. Branco, P., Torgo, L., Ribeiro, R.P.: A Survey of predictive modeling on imbalanced domains. ACM Comput. Surv. **49**(2), 31:1–31:50 (2016). doi:10.1145/2907070
4. Casner, S.M., Hutchins, E.L., Norman, D.: The challenges of partially automated driving. ACM Commun. **59**(5), 70–77 (2016). doi:10.1145/2830565
5. Davidoff, S., Ziebart, B.D., Zimmerman, J., Dey, A.K.: Learning patterns of pick-ups and drop-offs to support busy family coordination. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 2011), pp. 1175–1184 (2011). https://doi.org/10.1145/1978942.1979119
6. Dix, A.: Beyond intention-pushing boundaries with incidental interaction. In: Proceedings of Building Bridges: Interdisciplinary Context-Sensitive Computing (2002). https://pdfs.semanticscholar.org/d0a8/28c43654b1e16527a92863732cb13189fba9.pdf. Accessed 10 June 2017
7. Doryab, A., Min, J.K., Wiese, J., Zimmerman, J., Hong, J.I.: Detection of behavior change in people with depression. In: Proceedings of AAAI Workshop on Modern Artificial Intelligence for Health Analytics (MAIHA) (2014). http://repository.cmu.edu/hcii/271
8. Fogarty, J., Hudson, S.E., Atkeson, C.G., Avrahami, D., Forlizzi, J., Kiesler, S., Lee, J.C., Yang, J.: Predicting human interruptibility with sensors. ACM Trans. Comput. Hum. Interact. **12**(1), 119–146 (2005). doi:10.1145/1057237.1057243
9. Fogarty, J., Hudson, S.E., Lai, J.: Examining the robustness of sensor-based statistical models of human interruptibility. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 2004), pp. 207–214 (2004). https://doi.org/10.1145/985692.985719
10. Harrison, C., Tan, D., Morris, D.: Skinput: appropriating the body as an input surface. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 2010), pp. 453–462, (2010). https://doi.org/10.1145/1753326.1753394
11. Herlocker, J.L., Konstan, J.A., Riedl, J.: Explaining collaborative filtering recommendations, pp. 241–250 (2000). https://doi.org/10.1145/358916.358995
12. Höök, K.: Steps to take before intelligent user interfaces become real. Interact. Comput. **12** (4), 409–426 (2000). doi:10.1016/S0953-5438(99)00006-5
13. Wendy, J., Leifer, L.: The design of implicit interactions: making interactive systems less obnoxious. Des. Issues **24**(3), 72–84 (2008). doi:10.1162/desi.2008.24.3.72
14. Langheinrich, M.: Privacy by design—principles of privacy-aware ubiquitous systems. In: Ubicomp 2001: Ubiquitous Computing, pp. 273–291 (2001). https://doi.org/10.1007/3-540-45427-6_23
15. Langley, P.: Machine learning for adaptive user interfaces. In: Brewka, G., Habel, C., Nebel, B. (eds.) KI 1997. LNCS, vol. 1303, pp. 53–62. Springer, Heidelberg (1997). doi:10.1007/3540634932_3
16. Lim, C.H., Vats, E., Chan, C.S.: Fuzzy human motion analysis: a review. Pattern Recogn. **48** (5), 1773–1796 (2015). doi:10.1016/j.patcog.2014.11.016
17. McMillan, D., Loriette, A., Brown, B.: Repurposing conversation: experiments with the continuous speech stream. In: Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI 2015), pp. 3953–3962 (2015). https://doi.org/10.1145/2702123.2702532
18. Schmidt, A.: Implicit human computer interaction through context. Pers. Technol. **4**(2–3), 191–199 (2000). doi:10.1007/BF01324126
19. Smith, I.: Amazon releases Echo data in murder case, dropping first amendment argument. PBS (2017). http://www.pbs.org/newshour/rundown/amazon-releases-echo-data-murder-case-dropping-first-amendment-argument/
20. Wooldridge, M., Jennings, N.R.: Intelligent agents: theory and practice. Knowl. Eng. Rev. **10**(2), 115–152 (1995). doi:10.1017/S0269888900008122

21. Xiao, R., Laput, G., Zhang, Y., Harrison, C.: Deus EM Machina: on-touch contextual functionality for smart IoT appliances. In: Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI 2017), pp. 4000–4008 (2017). https://doi.org/10.1145/3025453.3025828
22. Yang, R., Newman, M.W.: Learning from a learning thermostat: lessons for intelligent systems for the home. In: Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp 2013), pp. 93–102 (2013). https://doi.org/10.1145/2493432.2493489

Social Media and Online Interaction

Open-Source Monitoring, Search and Analytics Over Social Media

Manos Schinas[1,2(✉)], Symeon Papadopoulos[1], Lazaros Apostolidis[1], Yiannis Kompatsiaris[1], and Pericles A. Mitkas[2]

[1] Centre for Research and Technology Hellas, Thessaloniki, Greece
{manosetro,papadop,laaposto,ikom}@iti.gr
[2] Aristotle University of Thessaloniki, Thessaloniki, Greece
mitkas@auth.gr

Abstract. The paper describes a technical demonstration of an open-source framework for monitoring, analysis and search over multiple social media platforms. The framework is intended to be a valuable tool for media intelligence professionals, as well as a framework and testbed for scientists and developers with interest in social media research.

Keywords: Social media · Data collection · Data mining · Information retrieval · Stream processing · Information visualization

1 Introduction

Social media [1] is nowadays a key channel for communication, self-expression, information gathering and entertainment for billions of Internet users worldwide. The penetration of social media has consistently increased throughout the last decade and it has nowadays reached a point where the large majority of people worldwide regularly use one or more of numerous popular platforms for a wide variety of purposes. As a result, social media is often regarded as a sensor of real-world trends and events [2] and as an indispensable tool for performing social science research [3] and consumer intelligence gathering and marketing [4].

Existing social media research is typically based on one of two approaches depending on the setting where it is carried out. In *academic* settings, researchers typically use a variety of tools, scripts or libraries, and often develop their own additional custom scripts to perform the required data collection, manipulation and analysis. In *business* settings, analysts use a variety of software-as-a-service products, typically through a dashboard-like interface, offering access to statistics and analytics about queries of interest. The main disadvantage of the first approach is the increased effort that is necessary to set up the data collection and analysis pipeline. In contrast, the second approach suffers from limited flexibility, cost (most social media SaaS offerings are subscription-based) and dependence on proprietary solutions.

To address these limitations, we present an integrated open-source framework for monitoring, analyzing and retrieving social media content. The framework is

I. Kompatsiaris et al. (Eds.): INSCI 2017, LNCS 10673, pp. 361–369, 2017.
https://doi.org/10.1007/978-3-319-70284-1_28

easy to install locally and can then be managed as a set of services that expose data collection, indexing and retrieval capabilities via a REST API and a web-based user interface. In addition, the framework is designed in a modular way, which makes it straightforward to adapt and extend for different requirements, e.g. collect data from additional social media platforms, filter incoming content, compute additional metrics, etc. To our knowledge, the proposed framework is the only integrated open-source solution for social media monitoring, analytics and search, which is freely available and offers at the same time ease-of-use, rich off-the-shelf features, flexibility and extensibility.

The development of the presented framework started in SocialSensor[1], where a need came up to retrieve multimedia content around trending topics that were automatically detected by the news monitoring application developed within the project [5]. The development of the tool continued in REVEAL[2], while the current version presented in this paper is based on extensions made to support the social media monitoring needs of the STEP[3] and hackAIR[4] research projects.

2 Design and Implementation

The presented framework is built upon a set of projects, all of which are available on GitHub[5]. The core project (*mklab-framework-common*) contains a set of classes for defining a common data model across different social media platforms (Fig. 1). Social media abstractions (*mklab-socialmedia-abstractions*) implements the data model for each of the supported social media platforms, e.g. Twitter, Facebook, etc., and contains a set of wrappers that encapsulate the different APIs provided by the platforms for the collection of data. The client project (*mklab-framework-client*) contains a set of classes for data management, including wrappers around different storage solutions, such as MongoDB and Apache Solr. Finally, the Stream Manager (*mklab-stream-manager*) incorporates all the orchestration and operational logic for data collection and storage.

2.1 Data Model

Figure 1 depicts the data model used to store data in the platform. The left side of the Entity-Relationship (ER) diagram depicts objects collected by the platform: *items*, *media items*, *web pages* and *users*. Items correspond to the messages posted to the various social media platforms, e.g., status updates on Twitter, posts on Facebook, video posts on YouTube, etc. Media items correspond to the multimedia content that is embedded in these items, for instance, images in tweets, videos in Facebook posts, etc. Web pages correspond to the URLs that are contained in the items. Finally, users are the objects representing the user

[1] http://socialsensor.eu/.
[2] https://revealproject.eu/.
[3] http://step4youth.eu/.
[4] http://www.hackair.eu/.
[5] https://github.com/MKLab-ITI/.

accounts publishing the items. The right side of the diagram depicts the *collection* entity, which offers users of the platform a way to organize social media content. A collection can comprise multiple queries as described in Sect. 2.2, while the same query can be included in more than one collections. The platform supports three types of query: (a) keyword queries, (b) account queries, and (c) location queries. It is important to note that there is no explicit (stored) association between a collection and the collected content (items, media items, etc.), nor is there such association between content and queries. This is depicted in the ER diagram by the dotted relationships between the corresponding entities.

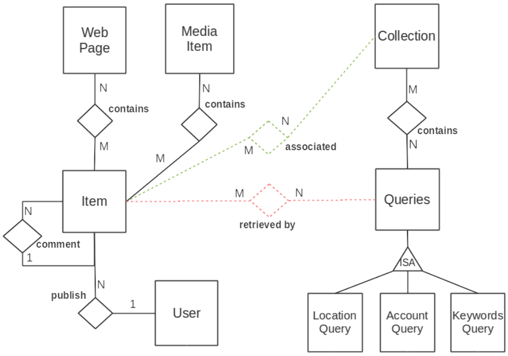

Fig. 1. Entity relationship data model of social media monitoring tool.

2.2 Data Collection and Processing

The collection of social media data is performed by the Stream Manager module. Currently, five different sources are supported: Twitter, Facebook, YouTube, Google+ and Flickr. The Stream Manager collects posts of interest (as defined by user queries) shared in these platforms alongside the users that published them, the embedded media items and the linked URLs. In addition to social media sources, the platform also supports monitoring RSS feeds.

Data collection by the Stream Manager is configured on the basis of user-defined collections, which, as mentioned in Sect. 2.1, consist of a set of three types of query: keywords, accounts, locations. Keywords can be defined on the basis of simple keywords, e.g. *"air pollution"*, or more complex logical expressions, e.g. *"election results" AND ("donald trump" OR "hilary clinton")*. Accounts refer to public sources of content for each platform, e.g., Twitter users, Facebook pages, YouTube channels, etc. Locations are represented as bounding boxes defined by pairs of latitude/longitude coordinates.

Given a collection, the associated queries are translated to the appropriate API calls. For example if a Twitter account such as *@BBC* is specified as one of the queries in the collection, that account is mapped to a call to the Twitter API that is used for the collection of Tweets posted by the specific user. In the same way, given a logical expression of keywords, multiple API calls are generated, one for each of the supported sources. However, given that the tool may be used by several end users, several of the defined collections may contain identical queries, e.g. the same keywords. To make the query process more efficient, the Stream Manager first de-duplicates all input queries into unique non-redundant query elements. Then, the Stream Manager periodically polls each of these unique query elements, and keeps track of the number of submitted requests to each platform in order to respect the limits imposed by it. It is worth noting that the set of Stream Manager queries may change dynamically (e.g., when a user creates or updates a collection), and such changes are efficiently communicated to the Stream Manager through a Redis[6] message broker instance.

The fetched items are then processed by a sequence of *filters* and *processors*, executed within the Stream Manager, before being stored and indexed. In terms of filtering, a set of heuristic rules is applied to keep only items of high quality. For example, items with limited text content, or with too many hashtags and URLs are treated as spam messages and therefore discarded. In terms of processing modules, three indicative processors are provided off-the-shelf by the Stream Manager: language detection, named entity extraction, and MinHash signature extraction from the text of the item. Finally, the collected content elements (items, media items, users and web pages) are stored in a MongoDB[7] instance.

2.3 Indexing and Retrieval

Data indexing is based on Apache Solr[8]. For each item, a subset of its fields are indexed. This includes all the textual fields, e.g. title, and other relevant fields, e.g., publication time, user id, number of views, etc., which can be used for filtering and faceting. To retrieve content related to a specific collection a Solr query is generated based on the queries that are associated with the collection, and the ids of the relevant items are retrieved. Then, the corresponding item metadata and all associated entities are retrieved from MongoDB by id.

The platform analytics capabilities range from simple metrics such as the number of items or unique users, to more complex ones such as *reach* (estimated cumulative audience for a set of items) and *endorsement* (estimated sum of "likes" for a set of items). Top users, locations, tags and named entities (based on frequency of appearance) are also provided. Finally, the platform supports the generation and visualization of timelines with different time granularities.

Analytics operations are implemented on top of two Solr components: Faceting and Stats. For example, for the generation of the top active users in the collection (in terms of number of posted items), we make a facet request in Solr, using the user id as the field to be treated as a facet. The corresponding response by Solr contains the top N users along with the number of items per user. In a similar manner, using the Stats component we can calculate field-oriented statistics, e.g. the average number of views or sum of shares per collection.

[6] https://redis.io/.

[7] https://www.mongodb.com/.

[8] http://lucene.apache.org/solr/.

Finally, the analytics framework leverages two more Solr components to make easier the exploration of collections: Result Clustering and Result Collapse. Clustering is used to automatically discover groups of similar search hits, i.e. groups of items related to a specific topic. These identified topics can be used to narrow down the set of items associated with a collection. Collapse, on the other hand, is used to group search results based on the value of a field. In our case, we use collapse based on the MinHash signature of items to support de-duplication of content, i.e. items with the same signature collapse into a single item.

Figure 2 depicts the key software components of the framework along with the function that each performs. Further details regarding their deployment are provided in Sect. 3.3.

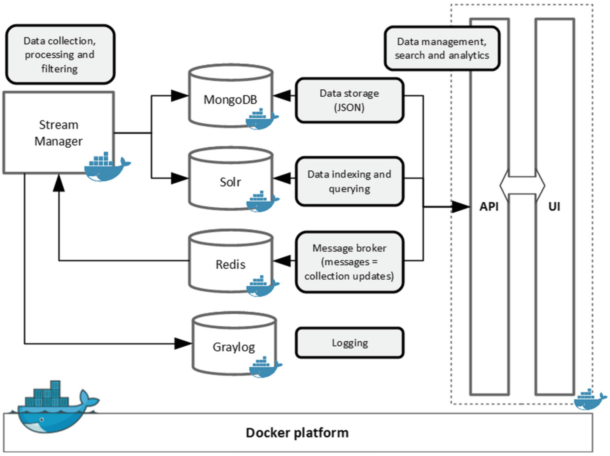

Fig. 2. Overview of framework components (white blocks) and their function (gray).

3 Usage and Deployment

3.1 Application Programming Interface

All the data collection, indexing, aggregation and retrieval capabilities of the framework are exposed to developers through a REST API that is built on PHP and served by an Apache web server. The methods of the API belong to three categories: (a) collection management, (b) content retrieval and (c) analytics. The first category contains a set of methods to create/update/delete collections and to retrieve collections created by a specific end user. For content retrieval there is the */items/collection-id* method, which is used to get the items related to a specific collection. The method has a set of parameters to refine how retrieval takes place, e.g. filtering or sorting. Finally, analytics over a selected collection is exposed through six methods: */statistics*, */timeline*, */users*, */terms*, */countries* and */heatmap/points*. For each of these methods the unique identifier of the collection must be specified. In addition, a set of parameters for the refinement of the collection (e.g. by source or language) is also provided.

3.2 User Interface

The framework user interface consists of three main pages: (a) a page where users of the platform can create their own collections, (b) a feed view for browsing content (Fig. 3), and (c) a dashboard view for analytics (Fig. 4).

The feed view presents the list of media items collected in the form of a stream. Each item in the view features relevant metadata, namely the publication time, username and social media platform. The feed can be presented in both gallery/grid (Fig. 3) and list form. The user can also search for items based on free text queries and use a number of filtering and sorting criteria. The following filters can be applied: (a) source (social media platform), (b) language, (c) topic, (d) original (show retweets/shares or not), (e) type (text item or item with embedded multimedia), (f) unique (remove near-duplicate content), (g) date. In addition to filtering, end users can rate the items retrieved for a collection on a scale of 1–5 (1: "irrelevant", 5: "relevant"). In addition, users may choose to exclude selected items and/or users along with their published items.

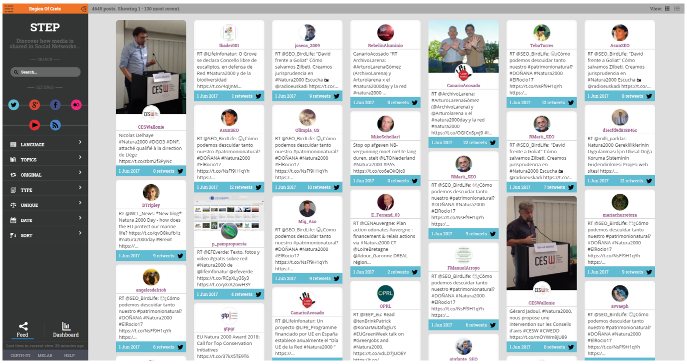

Fig. 3. Snapshot of feed gallery view.

The dashboard view offers metrics and widgets that depict summary views over the collected data. The visualizations are dynamic by leveraging the same set of filters as the ones in the feed view. In more detail, the dashboard consists of the following widgets:

- Numbers of posts made, users talking, users reached and endorsements, and platform contribution pie chart.
- Heatmap based on the exact location (latitude, longitude) of geo-tagged items in the collection (typically a small subset of all items).
- User location at the level of a country, which is automatically detected from the location text field in users' profiles.
- Timeline visualization based on the number of items over time. The granularity of the depiction may be set to hour, day or week.

Fig. 4. Snapshot of dashboard view.

– Top N users with most posts in the collection. N can be set between 10 and 200. Avatar, username and total number of posts for each user is presented along with the link to their social profile pages.
– Top N entities in terms of frequency. N can be set between 10 and 200. Entities are organized in three frequency categories (often, occasionally, seldom) and three types (person, tag, location).

3.3 Deployment to Local Environment

For deploying the framework as a locally administered service, we opted for the use of Docker[9]. Each of the framework components is deployed as a separate Docker container[10]. Figure 2 depicts all components and the communication between them. For MongoDB and Solr, the official Docker images are used. For the Stream Manager, we build an image that clones a stable release of the module from its GitHub repository and generates an executable jar using maven. The REST API and user interface is built and deployed as a separate Docker image on top of Apache web server. Finally, the signaling of collection-related events to MongoDB is handled by a Redis instance, while the aggregation and storage of logs by a Graylog instance[11], both deployed on top of Docker.

Using Docker Compose[12], we configure all components in a single YAML[13] file. This only needs limited edits to be adapted to users' local environment. The

[9] https://www.docker.com.
[10] Available on: https://github.com/MKLab-ITI/mmdemo-dockerized.
[11] https://www.graylog.org/.
[12] https://docs.docker.com/compose/.
[13] http://yaml.org/.

only requirement is to set the directory paths for MongoDB, Solr, etc. to the local file system. Optionally, service ports can be edited to values other than the default ones. In addition to the YAML configuration file, a separate configuration file associated with the Stream Manager needs to be edited to enable the use of social media APIs, by providing user credentials for each platform.

4 Limitations and Future Work

There are three main limitations in the usage of the presented framework. A first limitation is due to API usage restrictions imposed by the social media platforms. For example, Twitter restricts the number of allowed calls in a time window of 15 min[14]. Other platforms have similar restrictions but with different time granularities. As the number of collections and associated queries increases, it is necessary to decrease the request rate for each query. This creates a delay in the discovery of new content, especially for queries (keywords) that exhibit high activity. Moreover, a limitation pertains to the large scale of the collected content items, which may easily reach the order of tens of millions. As the association of content to collections is dynamically resolved based on the Solr indexing mechanism, the response time of the tool is adversely affected by the size of Solr index. This is more pronounced in the case of analytics, which are calculated on the fly using the faceting mechanism of Solr. A third limitation is the lack of built-in support in the framework for managing the privacy risks that are inherent when someone performs data collection from social media. In particular, the framework does not support secure data storage and access (e.g. via encryption), nor does it offer any facilities for removing user-generated content that could accidentally reveal personal information.

To improve the scalability of the framework, we plan to investigate and develop means of distributing the query and indexing load over multiple nodes (e.g., via sharding and replicating the index). In terms of functionality, we aim at two further improvements: (a) more sophisticated relevance models that go beyond the simple heuristic rules that are currently used for filtering, (b) support for additional social media platforms such as Sina Weibo and VKontakte.

Acknowledgements. This work was supported by the STEP (H2020-649493) and HackAir (H2020-688363) projects, funded by the European Commission.

References

1. Kaplan, A.M., Haenlein, M.: Users of the world, unite! the challenges and opportunities of social media. Bus. Horiz. **53**(1), 59–68 (2010)
2. Aiello, L.M., Petkos, G., Martin, C., Corney, D., Papadopoulos, S., Skraba, R., Goker, A., Kompatsiaris, Y., Jaimes, A.: Sensing trending topics in twitter. IEEE Trans. Multimedia **15**(6), 1268–1282 (2013)

[14] https://dev.twitter.com/rest/public/rate-limiting.

3. Mejova, Y., Weber, I., Macy, M.W.: Twitter: A Digital Socioscope. Cambridge University Press (2015)
4. Fan, W., Gordon, M.D.: The power of social media analytics. Commun. ACM **57**(6), 74–81 (2014)
5. Diplaris, S., Papadopoulos, S., Kompatsiaris, I., Goker, A., Macfarlane, A., Spangenberg, J., Hacid, H., Maknavicius, L., Klusch, M.: Socialsensor: sensing user generated input for improved media discovery and experience. In: Proceedings of the 21st International Conference on World Wide Web, WWW 2012 Companion, 243–246. ACM (2012)

"Reputational Heuristics" Violate Rationality: New Empirical Evidence in an Online Multiplayer Game

Mirko Duradoni[1,4](✉), Franco Bagnoli[2,4,5], and Andrea Guazzini[3,4]

[1] Department of Information Engineering, University of Florence, Florence, Italy
mirko.duradoni@unifi.it
[2] Department of Physics and Astronomy, University of Florence, Florence, Italy
franco.bagnoli@unifi.it
[3] Department of Education and Psychology, University of Florence, Florence, Italy
andrea.guazzini@unifi.it
[4] Center for the Study of Complex Dynamics, University of Florence, Florence, Italy
[5] INFN, sez.Firenze, Florence, Italy

Abstract. "Reputation systems" are widely used in e-commerce, crowdsourcing and crowdfunding platforms, as well as in a multitude of different web-based services. However, recent works stressed how the attribution of the reputation could be unrelated to the actual behaviour of the users. The aim of this study was to investigate which factors influenced the formation and the maintenance of the reputation in an online multiplayer game. Our study provided further and novel evidence of how people greatly rely on the previous acquired reputation of their interactors, whenever they are asked to rate them after a game's interaction. The "Reputational heuristics" adopted by players appeared to neglect the actual interactor's behaviour, in favor of a judgement in accordance with his behaviour.

Keywords: Reputation systems · Reputational heuristics · Human virtual dynamics · Online cognitive dynamics

1 Introduction

In virtual environments reputation based systems are nowadays very common and diffused in an increasing multitude of web based services. Information and Communication Technologiesn (ICTs) have facilitated the proliferation of systems based on online feedbacks [1]. For instance, e-commerce sites (e.g., e-Bay, Amazon) rely massively to them in order to discover the sellers' reliability, thus ensuring their own survival in the labour market. The same companies increasingly recognize the benefits associated with having a good reputation, and employ a considerable amount of resources and energy in reputation management practices [2–4]. Indeed, reputation seems able to exert an influence

I. Kompatsiaris et al. (Eds.): INSCI 2017, LNCS 10673, pp. 370–376, 2017.
https://doi.org/10.1007/978-3-319-70284-1_29

upon people behaviours, judgments, feelings and thoughts. Already Semmann et al. [5] noticed how people tended to be influenced greatly by the reputation of their interactors. Interestingly, also a huge set of projects concerning Collective Intelligence platforms owe their success to the ability of reputation systems to engage people and keep their interest alive through gamification [6]. For istance, the possibility to get badges, to climb charts or to acquire a certain level is what allowed Waze (i.e., a crowdsourced traffic monitoring application) to work and to be appreciated by a large number of users. The same results have been reached by applications related to education and learning goals, such as Duolingo and Khan Academy. However, despite all the aforementioned projects use reputation systems for their own well-functioning, barely new scientific evidence [7,8] stressed the fact that reputation could be acquired in a manner partially disconnected from the "actual trustworthiness" of our own social partners. Thus, it became a very important issue to understand what happens when reputation is derived from the feedbacks of other individuals (e.g., e-commerce sites) and not from people's actual behaviour as in some experimental settings [5]. Is well known how individuals often use cognitive heuristics in their social decision making process [9], and how such "algorithms" frequently appear to violate the principles of rationality [10,11]. For what concern the social judgments, people internalize behavioural rules of thumb from the very first years of life [12], and make their application automated. Moreover, as pointed out by Postmes and Spears [13], when people interact within a virtual environment, a psychological state defined as "de-individuation" could occur. In these circumstances the sensitivity of the individuals to local norms (i.e., those environmental signals that indicate which behaviour is appropriate and desirable in that context) increases. Moreover, the anonymity and the physical isolation that characterize many virtual interactions, have been found to push individuals to be more influenced by the set of local norms (e.g., subjects' reputation). Therefore, the aim of the present study was on the one hand to investigate whether or not individuals behaved as rational agents during reputational multiplayer online game (i.e., giving a positive feedback when they receive a gain, and providing a negative feedback when they receive a damage, disregarding the opponents' reputation). On the other hand, we wanted to make more clear which factors contribute to develop a good (or a bad) reputation during the virtual dynamics generated by the game. In order to investigate such a dynamics, we developed a Bargaining game with a reputation system, using the *Google Script* programming language. The game asked the participants to provide an evaluation of their opponents after certain game interactions, having clear their actual behaviour and current reputations.

2 Participants

The research was conducted respecting the guidelines for the ethical treatment of human participants of the Italian Psychological Association (i.e., AIP). A total of 113 subjects carried out entirely our experiment. 77 of them were adolescents (36 females) with an average age of 16 (s.d. 1.28) and were recruited in a high school in Prato (IT). While, 36 of them were adults (19 females) with an average age of 21 (s.d. 1.88). Both adolescents and adults were recruited through a complete voluntary census.

3 Methods and Procedures

Procedures and setting. The adolescents carried out the experiments in the computer lab of their high-school, while the adult participants took part in the experimental sessions in the computer lab of the Faculty of Psychology in Florence. Despite these two samples accomplished the experiments in two different locations the procedures and the lab environment were made as identical as possible. The participants seated at their computer positions that were separated by the others by means of partitions. Moreover, the subjects were strictly not allowed to talk one another and all the interaction between them were made anonymously via computer. Before receiving their access data, the instructions of the game were showed on the participants' monitor and read aloud by the experimenter as well.

Bargaining game. The game involved group of six players interacting one another for 45 rounds (15 in each role). In this game Receivers had to decide to accept or to decline a deal (exchange of resources) proposed by another player (i.e., Donor). The Receivers could only saw the amount and type of the resource offered by the Donor, but were unaware of what and how much the Donor asked them in return. To take their decision the Receivers could asked for a suggestion to an Observer that was identifiable only from his reputation. Indeed, the Receivers' available information about the Observers was limited to their reputation score. Thus, the Receivers were fully unaware about the actual Observer's previous moves. The Receivers couldn't select the Observer they wanted from all the other players that were playing that role, but they were matched each turn with a random Observer by our system. At the end of the game each Receiver was matched exactly three times with each Observer. However, matching did not ensure that the Receivers interacted (i.e., ask for a suggestion) with the matched Observer. Once the Receiver's decision has been taken, the Donor's request is revealed. If the Receiver asked for the Observer's suggestion (to accept or to decline the Donor's offer), he had the opportunity to feedback the Observer and thus contribute to the Observer's reputation. We specify that the players could rate the observer only after becoming aware of the real request of the Donor. In this way the receivers always knew if the observer had been "fair" with them or not.

Data Analysis
The preconditions necessary for the inferential analyzes were verified on the experiment's data. For the continuous observables that were under investigation, the normality of the distribution was assessed through the analysis of asymmetry and kurtosis values. Due to the repeated measures structure of the experimental data, the inferential analyses were conducted using a general linear mixed model (GLMM).

4 Results

To understand how reputation was attributed within our setting we focused our attention to the feedback behaviour of the Receivers. Indeed, was through these actions (give a like or a dislike) that the Observer's reputation was built. Therefore, we analysed the feedback behaviour by means of some generalized linear mixed models. The gender and the age of the participants as well as the game-related observables (i.e., the goodness of the suggestion provided, the previously acquired reputation) were all considered as parameters.

For what concerned the game-related variables, we defined them as it follows:

- The Observer provided a good suggestion when he suggested to refuse an offer with a negative difference between the amount offered and the amount required in return, and also when he suggested to accept a positive exchange offer. Instead, suggest to the Receiver to accept a disadvantageous deal as well as to refuse a positive one were classified as bad observations. A total of 810 suggestions have been categorized as good suggestion while 421 as bad suggestions.
- The Observer's reputation was defined by the difference between the likes and the dislikes received. Thus, a positive reputation was characterized by a positive difference between these two feedbacks while a bad one by a greater number of dislikes. In total 543 game records referred to negative reputational scores while 610 to good rated Observers.

To better represent the interaction between these two game-related variables and their distribution in our sample, in Table 1 we presented the percentage of good and bad suggestions for each reputation category (i.e., positive and negative).

Table 1. Good and bad suggestions across reputational scores

	Bad suggestion	Good suggestion
Negative reputation	36.0%	64.0%
Positive reputation	30.5%	69.5%

The final and best model is reported in Table 2.

The Age and the Gender of the participants did not seem to influence the feedback behaviour of the Receiver neither directly nor trough interaction effects. Only two factor appeared to contribute to define the Observers' reputation. Specifically, good suggestions and positive reputations drew more frequently a positive feedback from the others. Conversely, providing a bad observation or having a bad reputational score, determined less positive feedbacks and more negative ones. Interestingly, the standardized β seems to highlight how Receivers could have been more affected in their feedback decision making by the previous

Table 2. Generalized linear mixed models. Factors that influence the feedback behaviour of the receivers.

GLMM best model LIKE				
	Model precision	*Akaike**	F	Df-1(2)
Best model	$74,4\%$	24.515	42.265^{***}	2(317)
Fixed effects				
Factor		**F**		**Df-1(2)**
Reputation		74.154^{***}		1(317)
Goodness of suggestion		27.078^{***}		1(317)
Parameter		**Coefficient (β)**		**Student t**
Reputation(-)		-2.257		-8.611^{***}
Goodness of suggestion(-)		-1.448		-5.203^{***}

$^{***} = p. < 0.001$; Reputation (-): Bad reputation; Goodness of suggestion (-): Bad suggestion.

Table 3. Generalized linear mixed models. Factors that influence the feedback behaviour of the receivers when they have received a good suggestion from the observer.

GLMM best model LIKE - Good suggestion received				
	Model precision	*Akaike**	**F**	**Df-1(2)**
Best model	$76,5\%$	13.422	49.251^{***}	1(215)
Fixed effects				
Factor		**F**		**Df-1(2)**
Reputation		49.251^{***}		1(215)
Parameter		**Coefficient (β)**		**Student t**
Reputation(-)		-2.443		-7.018^{***}

$^{***} = p. < 0.001$; Reputation (-): Bad reputation.

acquired reputation of their interactor instead of his actual behaviour (i.e., the goodness of the suggestion received).

The influence exerted by the reputation on the Receivers' judgment has been further investigated through two other generalized linear mixed models. In each one of them we selected only one typology of suggestion. The first model (Table 3) refers to those cases in which the Receivers obtained good suggestions from the Observers, while the second one (Table 4) considered only those situations in which the Observers provided bad observations to the Receivers.

When the Observers provided a good suggestion, those of them with a good reputation attracted more positive feedback from the others respect to those with a negative reputational score (Table 3). The same pattern has been observed in those cases in which the Observers' suggestion damaged the Receivers (Table 4). Overall, for equal suggestion received (i.e., identical actual behaviour) we observed a preference for those Observers who gained previously a good reputation.

Table 4. Generalized linear mixed models. Factors that influence the feedback behaviour of the receivers when they have received a bad suggestion from the observer.

GLMM best model LIKE - Bad suggestion received				
	Model precision	*Akaike**	F	Df-1(2)
Best model	$70,9\%$	12.325	16.611^{***}	1(101)
Fixed effects				
Factor		**F**		**Df-1(2)**
Reputation		16.611^{***}		1(101)
Parameter		**Coefficient (β)**		**Student *t***
Reputation(-)		-1.897		-4.076^{***}

$^{***} = p. < 0.001$; Reputation (-): Bad reputation.

5 Discussion

Understand how reputation is constructed within a widespread feedback system, like the ones used in e-commerce as well as in some crowdsourcing projects, is a major issue. Indeed, reputation systems are actually considered as the most effective mechanism to foster cooperation in virtual environments [14]. However, recent works suggested how this process could be biased [7,8]. Our work provided novel evidence about how reputation is really "made" within such virtual environments. Differently from rational agents, people are not only influenced in their judgments by the actual behaviour of their interactor. Indeed, they seemed to be greatly affected also by the previous acquired reputation of their social partner. In general, people preferentially rewarded (i.e., provided a positive feedback) individuals who gained in the past interactions a good reputation and conversely punished (i.e., provided a negative feedback) more often participants who obtained a bad reputational score. This feedback tendency is maintained disregarding the actual behaviour of the social partner. Indeed, equally trustworthy individuals (i.e., participants which shown the same suggestion behaviour) are treated (i.e., feedbacked) differentially according to their reputational rating. To put it simply, people seemed to use a sort of reputational "heuristics" to take a feedback decision that violated rationality (i.e., appeared to be disconnected from the actual behaviour). Our results highlighted a critical aspect of the use of reputational systems in virtual environments. Indeed, given this feedback tendency, reputation appear to preserve its state (i.e., to be maintained) even despite the actual behaviour of the social partner. In other words, reputation appear to have an inertia. Not only that, initial distinctions in reputation levels leaded to a different reputation rate grow. Indeed, good-rated Observers acquired more easily further social rewards whether they helped or damaged the Receivers respect to those Observers who were identified as bad partners. The great influence exerted by reputation within a widespread feedback system appear in line with de-individuation studies [13]. Indeed, reputation in such environments could convey that important social information to which people are very sensitive. The anonymity and the physical isolation that characterized such environments appear able to trigger the reputation heuristics decision making. This process ends up in reinforcing that social evaluations (i.e., reputation) that has been constructed by the whole group through the widespread feedback system. In the end our results stressed the necessity to further investigated the circumstances under which the reputation systems could work properly.

Acknowledgement. This work was supported by EU Commission (FP7-ICT-2013-10) Proposal No. 611299 SciCafe 2.0.

References

1. Dellarocas, C.: The digitization of word of mouth: promise and challenges of online feedback mechanisms. Manag. Sci. **49**(10), 1407–1424 (2003)
2. Dolle, R.: Online reputation management (Bachelor's thesis, University of Twente) (2014)
3. Hutton, J.G., Goodman, M.B., Alexander, J.B., Genest, C.M.: Reputation management: the new face of corporate public relations? Publ. Relat. Rev. **27**(3), 247–261 (2001)
4. Waeraas, A., Byrkjeflot, H.: Public sector organizations and reputation management: five problems. Int. Publ. Manag. J. **15**(2), 186–206 (2012)
5. Semmann, D., Krambeck, H.J., Milinski, M.: Reputation is valuable within and outside ones own social group. Behav. Ecol.Sociobiol. **57**(6), 611–616 (2005)
6. Katmada, A., Satsiou, A., Kompatsiaris, I.: Incentive mechanisms for crowdsourcing platforms. In: Bagnoli, F., Satsiou, A., Stavrakakis, I., Nesi, P., Pacini, G., Welp, Y., Tiropanis, T., DiFranzo, D. (eds.) INSCI 2016. LNCS, vol. 9934, pp. 3–18. Springer, Cham (2016). doi:10.1007/978-3-319-45982-0_1
7. Duradoni, M., Bagnoli, F., Guazzini, A.: Non-trivial reputation effects on social decision making in virtual environment. In: Bagnoli, F., Satsiou, A., Stavrakakis, I., Nesi, P., Pacini, G., Welp, Y., Tiropanis, T., DiFranzo, D. (eds.) INSCI 2016. LNCS, vol. 9934, pp. 115–122. Springer, Cham (2016). doi:10.1007/978-3-319-45982-0_11
8. Frey, V., van de Rijt, A.: Arbitrary inequality in reputation systems. Sci. Rep. **6**, 38304 (2016)
9. Rand, D.G., Peysakhovich, A., Kraft-Todd, G.T., Newman, G.E., Wurzbacher, O., Nowak, M.A., Greene, J.D.: Social heuristics shape intuitive cooperation. Nat. Commun. **5**, 3677 (2014)
10. Kahneman, D., Tversky, A.: Prospect theory: an analysis of decision under risk. Econometrica: J. Econometric Soc. **47**(2), 263–291 (1979)
11. Fehr, E., Gachter, S.: Altruistic punishment in humans. Nature **415**(6868), 137 (2002)
12. Olson, K.R., Spelke, E.S.: Foundations of cooperation in young children. Cognition **108**(1), 222–231 (2008)
13. Postmes, T., Spears, R.: Deindividuation and antinormative behavior: a meta-analysis. Psychol. Bull. **123**(3), 238 (1998)
14. Wu, J., Balliet, D., Van Lange, P.A.: Gossip versus punishment: the efficiency of reputation to promote and maintain cooperation. Sci. Rep. **6**, 23919 (2016)

Why People Use Chatbots

Petter Bae Brandtzaeg(✉) and Asbjørn Følstad

SINTEF, Forskningsveien 1, 0314 Oslo, Norway
petter.b.brandtzag@sintef.no

Abstract. There is a growing interest in chatbots, which are machine agents serving as natural language user interfaces for data and service providers. However, no studies have empirically investigated people's motivations for using chatbots. In this study, an online questionnaire asked chatbot users (N = 146, aged 16–55 years) from the US to report their reasons for using chatbots. The study identifies key motivational factors driving chatbot use. The most frequently reported motivational factor is "productivity"; chatbots help users to obtain timely and efficient assistance or information. Chatbot users also reported motivations pertaining to entertainment, social and relational factors, and curiosity about what they view as a novel phenomenon. The findings are discussed in terms of the uses and gratifications theory, and they provide insight into why people choose to interact with automated agents online. The findings can help developers facilitate better human–chatbot interaction experiences in the future. Possible design guidelines are suggested, reflecting different chatbot user motivations.

Keywords: Chatbots · Motivations · Uses and gratifications

1 Introduction

Chatbots represent a potential shift in how people interact with data and services online. While there is currently a surge of interest in chatbot design and development, we lack knowledge about why people use chatbots.

Chatbots are machine agents that serve as natural language user interfaces for data and service providers [1]. Currently, chatbots are typically designed and developed for mobile messaging applications [2].

The current interest in chatbots is spurred by recent developments in artificial intelligence (AI) and machine learning. Major Internet companies such as Google, Facebook, and Microsoft see chatbots as the next popular technology; Microsoft CEO Satya Nadella said, "Chatbots are the new apps" [3]. In Spring 2016, Facebook and Microsoft provided resources for creating chatbots to be integrated into their respective messaging platforms, Messenger and Skype. One year later, more than 30,000 chatbots have been launched on Facebook Messenger. Other messaging platforms have also seen a substantial increase in chatbots, including Slack, Kik, and Viber. Chatbots are seen as a means for direct user or customer engagement through text messaging for customer service or marketing purposes [4], bypassing the need for special-purpose apps or webpages.

I. Kompatsiaris et al. (Eds.): INSCI 2017, LNCS 10673, pp. 377–392, 2017.
https://doi.org/10.1007/978-3-319-70284-1_30

However, it is not simple to transition from established user interfaces, such as web pages and apps, to chatbots as a common means of interacting with data and services. For example, there is a lack of knowledge regarding how customers react to the substitution of human customer service personnel with chatbots or how the presence of chatbots in online social networks affects multi-party conversations and the spread of information [5].

Since the initial optimism regarding the launch of chatbots by Microsoft and Facebook, a number of commentators have noted that users' adoption of available chatbots is less substantial than hoped [6]. This could be explained by the fact that most available chatbots fail to fill users' needs due to unclear purposes, nonsensical responses, or insufficient usability [7].

Designing a new interactive technology such as a chatbot requires in-depth knowledge of users' motivations for using the technology, which allows the designer to overcome challenges regarding the adoption of the technology [8]. More general knowledge is also needed to understand human–chatbot relationships. To our knowledge, no studies to date have investigated users' motivations for interacting with chatbots.

As a first step towards bridging this knowledge gap, we perform a study addressing the following research question:

RQ: *Why do people use chatbots?*

The study contributes new knowledge regarding individuals' motivations for using chatbots based on an online questionnaire completed by US chatbot users. The questionnaire includes an open question regarding the participants' main motivations for using chatbots. The findings obtained using this approach can inform future designs intended to improve human–chatbot interactions.

Before we present the findings, we will first describe the relevant background for our study. We then present the method and findings of the study. In the discussion, we address the implications of the study's findings for the design and development of chatbots.

2 Background

2.1 Chatbots and Natural Language User Interfaces

Although the last few years have seen increased interest in chatbots, natural language interfaces are not new in the fields of computer science and Internet studies. In the 1960s, Weizenbaum published an innovative study on natural language interaction with ELIZA, a computer program developed to mimic the responses of a psychotherapist in a therapy session [9].

Dale discusses "the return of chatbots" in recent years and discusses how the current interest in this technology is rooted in previous work on natural language user interfaces [1]. In particular, Dale notes the impact of the Loebner Prize, which has driven natural language user interfaces to be more human-like since 1991, and Pandorabots, a chatbot platform that includes more than 200,000 bot developers as of 2016. The best-known chatbots are Cleverbot, which was publicly launched in 1997;

A.L.I.C.E., the winner of the Loebner Prize in 2000, 2001, and 2004; and Mitsuku, the winner of the Loebner Prize in 2013 and 2016.

The current interest in chatbots is likely related to substantial advances in computing technology and the wide adoption of mobile messaging applications.

First, recent advances in artificial intelligence and machine learning have led to the recent interest in chatbots. These advances promise vast improvements in natural language interpretation and prediction capabilities, including improvements in machine translation [10]. In addition, progress in conversational modeling suggests that predictions based on recurrent neural networks and sequence-to-sequence models will outperform the rule-based conversational modeling typically applied to traditional chatbots [11].

Second, the increased adoption of mobile Internet and messaging platforms have driven the adoption of chatbots [2]. Through mobile messaging platforms, chatbots are able to reach a large part of the online population. According to *Business Insider* (2016), about 3 billion people worldwide use mobile messaging applications such as Facebook Messenger, WeChat, Skype, Telegram, Slack, Viber, and Kik. For many users of these services, natural language is expected in online interactions, making automated marketing and customer service using natural language a promising business opportunity.

2.2 Chatbot Applications

Chatbots may serve a number of purposes, such as customer service, social and emotional support, information, entertainment, and ties the user to other people or machines. The great variety of chatbots is exemplified in the BotList (https://botlist.co/), a website on which people can find chatbots for a broad range of purposes available on multiple messaging platforms.

In particular, chatbots are seen as a promising alternative to traditional customer service [4]. For customers, conversations with these bots may feel more natural and efficient than interacting with a mobile app as they can obtain answers to questions, receive suggestions for purchases, place orders, and keep updated on shipping through a natural language interface.

A range of chatbots serve as virtual assistants or stewards, helping users to perform specific tasks. The Indian chatbot Nikibot can help users with, for example, booking a taxi and ordering food for delivery [12]. In addition, Do Not Pay, based in the UK, helps users file complaints when they receive parking tickets and Babylon Health's chatbot interface provides medical advice. In such scenarios, chatbots may be preferable to other means of assistance, such as a phone call or online search, due to their convenience and immediacy.

Chatbots can also help people explore online content or services. For example, Microsoft launched Heston Bot to help users explore food and cooking opportunities Additionally, the global fashion and clothing company H&M launched a chatbot to provide personal fashion advice based on photos uploaded by users [13].

"Smalltalk" orientated chatbots such as Mitsuku and Jessie Humani can also fulfill people's need for entertainment and social interaction.

2.3 User Behavior and Experience

Although little is known about what motivates people to use chatbots, there is a substantial body of research on users' behavior and experience with chatbots.

Users' interactions with chatbots often mimic interactions between humans, but there are differences. In a study comparing human–human interactions to human–chatbot interactions, Hill et al. found that human–chatbot interactions tend to last longer than human–human interactions between strangers and involve shorter messages, less complicated vocabulary, and more profanity [14].

Corti and Gillespie investigated whether users seek to repair misunderstandings in conversations with natural language user interfaces, which is important in any type of dialogue [15]. They found that, for chatbots perceived as human, users made more of an effort to repair misunderstandings than did users that perceived the chatbots as automated [15].

Several studies have investigated users' experiences with chatbots. For example, Holtgraves et al. explored how users perceive chatbots' personalities [16], and De Angeli et al. studied how the implied anthropomorphism of chatbots may elicit negative responses among users [17]. Comparing the conversational ability of a chatbot based on the original ELIZA program to newer chatbots, a study [10] found that the participants were able to systematically differentiate the conversational quality of different chatbots. In addition, different demographic groups tended to rate the chatbots' conversational quality differently; specifically, younger users and female users rated the conversations more favorably [10].

It may be important for chatbots to engage emotionally with users. A recent study by Xu et al. on customer service chatbots found that about 40% of user requests to customer service are emotional rather than seeking specific information [4]. Without the ability to relate to these customers emotionally, a customer service chatbot risks failure.

2.4 Uses and Gratifications – A Theoretical Framework

As a theoretical basis for understanding people's motivation for using chatbots, we apply the well-established uses and gratifications theory (U&G) [18]. U&G explains why and how people use specific media to fulfill specific needs; the specific use of a medium depends on the expected and experienced gratification it will provide. U&G has typically been oriented towards consumers' use of media that is not related to work [19]. The theory assumes that the user is goal-driven in his or her selection and use of a particular medium based on social and psychological needs or gratifications.

In a complex media landscape or so called high-choice media environments, where users can choose to achieve their goals through a number of different media, such as webpages, apps, and chatbots, U&G assumes that the user takes an active stance towards which medium that best suits the purpose. Rubin describes U&G as a highly compatible approach to understanding the uses and effects of electronic media in the current media landscape [18].

The framework classifies users based on the gratification (or motivation) received from a particular medium, and it assumes that media users actively choose a medium depending on what they see as fit to satisfy a particular need [20]. Specifically, "uses and gratifications" refer to the motivation for use of a specific medium and the satisfaction people gain from use [21]. A wide range of gratifications have been suggested as motivators of media use [22], such as the need for information, entertainment, social interaction, and self-expression.

While the fragmentation of the media landscape due to the adoption of mobile Internet may change users' motivations for choosing certain media, Sundar and Limperos concluded that the gratifications for use of Internet technologies are similar to the gratifications for use of other media [22]. However, there are substantial variations between media contexts, which means that it is important to identify gratifications that are relevant to the context of the medium of interest. Identifying the gratifications that are important to chatbot users will help guide the development and design of new and existing chatbots.

Chatbots are a new technology, and as such, are mostly used by innovators and early adopters. These users might have different needs and gratifications than the rest of the population. The theory of diffusion of innovations explains how such innovations are adopted by a population. One of the insights that might be useful in combinations with U&G theory is the understanding the various user needs and gratifications among different user segments in the population, suggested by Rogers: Innovators: 2.5% Early Adopters: 13.5%, Early majority: 34%, Late majority 34%, Laggards 16%. Early adaptors are usually more risk-oriented and curious about new technologies, while the early and late majorities and laggards are more conservative and risk-averse [23].

3 Method

3.1 Study Design and Materials

To explore why people use chatbots and reach a sufficiently broad sample of chatbot users, a questionnaire was used. The questionnaire included 17 questions regarding chatbot use, including motivations for and experiences with chatbots, and respondents' demographics (age, gender, and state of residence).

In line with the exploratory aim of the study, the question addressing respondents' motivations for chatbot use was open: *What is your main reason for using chatbots?* The participants were asked to answer this open question freely. The question was adapted from a study by Brandtzæg and Heim [24] focusing on motivations for using social networking sites, as it encouraged participants in that study to provide personal descriptions of their motivations for using certain media.

Other key questions regarding chatbot use explored how often participants used chatbots, how long they had used chatbots, and the messaging platforms on which they used chatbots.

3.2 Participant Recruitment and Filtering

To understand users' motivations for chatbot usage, our target group in this study consists of only chatbot users. However, due to the newness of chatbots, most people in the mass market likely have not had any experience with chatbots, and many may not even know what chatbots are. Furthermore, no statistics regarding chatbot usage are available globally or for specific countries. Hence, recruiting participants for this study was challenging, as we had to not only identify relevant participants but also filter out non-relevant participants.

We decided to target chatbot users in the US as the technology companies that prioritize chatbots—Google, Facebook, Kik, and Slack—are all focused on the US market. We also decided to target a relatively young user group (those aged 16–55 years). We consider users within this age group to be more likely to be early adopters of chatbots than older users (e.g. [23]) as the former more frequently use messaging applications.

Data were collected in April 2017 by Survata, an independent US-based research company. Survata collect research samples by partnering with online publishers, which allow visitors to take a Survata survey to unlock premium content (e.g., premium articles, e-books, and videos). To identify and avoid invalid survey responses, Survata's technology analyzes respondents' response time, response pattern, and other metadata.

To identify relevant participants, the following screening question was applied: *Chatbots are automated online services that you interact with in text-based conversations, typically in instant messaging platforms such as Facebook Messenger, Kik, Slack, and Telegram. Have you used such chatbots?* Only those who responded positively to this question were allowed to take the survey.

As the screening question may not have been sufficient to filter out all non-chatbot users, we also analyzed users' responses to the open question. Specifically, we scrutinized the answers for indication that the participants discussed general use of messaging platforms rather than interaction with chatbots. This process identified 155 of the 301 participants as non-chatbot users. We are confident that the remaining 146 participants are actually within the target group of chatbot users.

3.3 Data Analysis

Data obtained from responses to the preset questions were analyzed through descriptive statistics using the SPSS 24 statistical package.

Qualitative data regarding participants' open answers underwent content analysis based on the coding categories established through an initial thematic analysis. Content analysis has proven to be useful for describing and making inferences about respondents' communications and patterns of usage as well as the consequences of communication [24].

The two authors collaboratively coded the open answer data in order to develop and apply the categories of motivation. Participants' responses to the open question could include more than one such category. These categories were then used by one of the authors to code the entire data set. In total, 16% of the responses were coded as addressing two or more themes. To ensure reliability, the other author reviewed the coding and made suggestions when necessary. The suggestions were reviewed by both authors, and then the final coding was performed. In total, 21% of the initial codes were updated.

4 Results

In total, 146 valid responses were gathered. Of these, 94 were written by females, and 52 were written by males. The mean age of the participants was 30 years (min = 16, max = 55, SD = 9.2). As shown in Fig. 1, the participants reported use of chatbots on various platforms, with Facebook Messenger being the most common, in line with the broad adoption of this platform in the US.

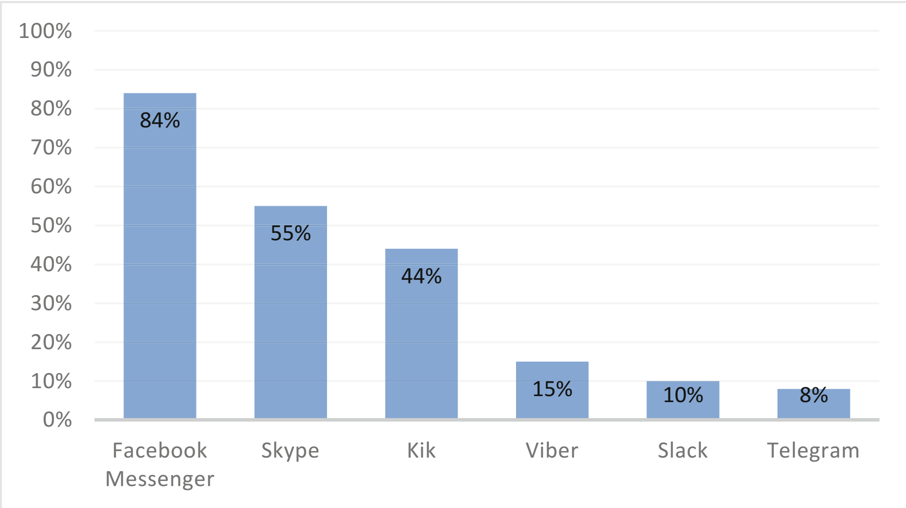

Fig. 1. Proportion of participants reporting use of different messaging platforms for chatbot interactions (N = 146).

In general, the participants were fairly new to using chatbots; 64% reported using chatbots for two years or less. Thirty-four percent of the participants reported using Google Assistant, a chatbot assistant available for select Android operating systems as well as on the Allo messaging platform.

In the open question regarding participants' motivations for using chatbots, all participants provided an answer. The thoroughness of the answers varied substantially, with the longest answer being 390 characters and the shortest being 5 (mean = 76; SD = 66).

The thematic analysis of the answers revealed four main categories of motivation. An overview of the categories is provided in Table 1. In the remainder of the results section, we present the detailed findings for each category.

4.1 Chatbots for Productivity

The vast majority of participants (68%) reported productivity to be the main reason for using chatbots. These participants highlighted the ease, speed, and convenience of using chatbots. Also, they noted that chatbots provide assistance and access to information.

Table 1. Categories of motivation for chatbot use (N = 146). Note: 16% of the responses were coded as addressing two or more themes.

Category	Description	Frequency
Productivity	The comment concerns the convenience of using chatbots (whether they are easy or fast to use). Participants typically report using them to obtain assistance or information	100
Entertainment	The comment concerns the entertainment value of using chatbots (whether they are fun to use). Some report that they use chatbots when bored to kill time	29
Social/relational	The comment concerns the use of chatbots for social or relational purposes. Typically, chatbots are seen as a personal, human means of interaction that may have social value. Some also use chatbots to strengthen social interactions with other people	18
Novelty/Curiosity	The comment concerns the use of chatbots out of curiosity or because they are a novelty. Often, the stated aim is to investigate chatbots' capabilities	15
Other	The comment concerns motivations that do not fit in the above categories and are not sufficiently frequent to justify a separate category	12

Ease, speed, and convenience

Forty-two percent of participants reported ease of use, speed, and convenience as their main reasons for using chatbots. These motivations are exemplified in the following statements:

> *To get answers quickly. It's important.* (P40)

> *It saves me the hassle of having to place a call, have to wait to speak to a person, then trying to get the information from that person. It also saves time in having to look through tons of text to find answers.* (P58)

> *They are fast and almost as fast as searching on the internet.* (P252)

As reflected in the above examples, the participants' responses suggested that they highly appreciated when chatbots helped them save time or made it easier and faster to obtain help or information, such as by providing efficient assistance in a customer support situation or by point to an easy-to-use manual or FAQ. Five percent of the participants explicitly stated that they use chatbots to avoid waiting for assistance, such as the following participant:

> *There is no wait time to talk to a representative to find out basic info I'm looking for. The chatbot can answer basic questions and is ready whenever I need an answer.* (P83)

Another five percent reported that they preferred to obtain help from a chatbot rather than a human assistant, such as in the following example:

Chatbots are important because you won't feel stupid asking important questions. Sometimes talking to someone can be a bit intimidating. Talking to a chatbot makes that a lot easier! (P82)

To obtain help and information
Many participants (41%) also reported that the ease of obtaining help or accessing information were main motivations for using chatbots. These people perceived chatbots as useful in their daily online activities, as exemplified by the following statements:

I use chatbots instead of a search engine to help with daily tasks. (P67)

I like to travel and find many different places to visits and chatbots gives me great advise. (P113)

Finding information about weather and news stories. (P251)

Eighteen percent of participants reported ease, speed, and convenience as well as help and information as motivation for using chatbots, including the following participant:

I believe the main reason for using chatbots is to give you help if you have any questions anytime when you need it, because the thing is when you ask a friend for help on kik or any time of messaging app they won't reply as quick as chatbots do. (P151)

Three participants reported that they appreciated being able to configure chatbots to receive better help with their personal tasks, as in the following statement:

I use chatbots because I can tailor them to find things, styles, weather, or orders that are specific and unique to me, so I can shop, research, or just chat with ease. The more input I put into them about me, the more they will serve efficiently. (P176)

4.2 Chatbots for Entertainment

A substantial proportion of participants (20%) reported using chatbots for entertainment. Though much smaller than the proportion of participants that use chatbots for increased productivity, this motivation category was the second most frequent.

Most of the participants reporting use of chatbots for entertainment value (14% of the total sample) described this as a positive value, with chatbots perceived as "fun" and "entertaining":

It's fun and entertaining. I like chatbots that have funny things to say. (P99)

Usually to ask a question and be entertained with an answer. (P301)

Others reported using chatbots for entertainment in a more negative fashion, to kill time:

I usually use chatbots when I am bored or have nothing to do so I use them to waste some time. (P61)

I am bored and want to talk to someone. (P232)

About one-third of the participants within this category (7% of the total participant sample) reported both productivity and entertainment as motivations to use chatbots, as reflected in the following statement:

> *Chatbots can help me with simple tasks on the app I'm using. Also, they give me fun tips and make my experience a lot better. It gives me something different to do when I'm bored.* (P199)

4.3 Chatbots for Social and Relational Purposes

The third most frequently reported reason for using chatbots is the potential social and relational benefits they can provide. This category of motivation was reported by 12% of participants.

It is noteworthy that, while chatbots can enhance interactions between humans, most of the participants addressing social and relational motivations commented on the social experience of interacting with the chatbot (10% of the total participant sample). For example, the chatbot is perceived as a way to avoid loneliness or fulfill a desire for socialization:

> *At the time i was bored and i didnt have anyone to talk to and i feel like sometimes their good to make friends with if your lonely and just want a chat with someone else.* (P141)

> *I use them when I feel bored or rather when I feel down and have no one else to go to, it just relaxes me in a way. Gives me someone to vent to without getting judged, I know they aren't real but it feels like it is.* (P264)

> *Always open to talk and I live in country and not many people around me to talk too.* (P234)

Some participants (3% of the total sample) reported that chatbots enhanced their social experiences with others, such as when including a chatbot in a group chat, using a chatbot with a child, or to improve one's own conversational skills:

> *To have a little extra fun in the chats I'm in.* (P268)

> *Normally to get information, but I also like using them for my kids to talk to characters.* (P60)

> *To build conversation skill.* (P80)

More than one-third of the participants reporting social and relational motivations (5% of the total participant sample) also mentioned productivity. The ability for chatbots to meet one's social and relational needs and improve productivity is seen as a benefit, as in the examples below:

> *So I can have someone to talk to AND it searches for me. Without me using the net.* (P243)

> *They are like personal assistants and easy to use because they are built into the apps themselves. I don't need to download extra apps I just talk to them directly on the app that I'm using, i.e. Facebook Messenger. I like that because it's a hassle using so many different apps. Also there's a sense of talking to someone when I use them. It's almost like you are talking to a real person.* (P116)

4.4 The Novelty of Chatbots

The fourth main category of motivation for using chatbots is the novelty of chatbots, reported by 10% of participants. These participants typically stated that they are curious to explore chatbots and the limits of their abilities:

> *Trying something new.* (P104)

> […] *I'm also curious to see what they'll say or how realistic they seem.* (P59)

> […] *It's interesting to see what people can come up with, how lifelike they will become. Sadly, very few pass the test. They are all repetitive in some way.* (P88)

Some of these participants seem to be attracted to the fact that chatbots are still in an early phase of development, suggesting that they enjoy being early adopters of technology:

> *They're new and intriguing.* (P66)

Others were skeptical of the new interactive technology:

> *Curiosity, mostly, because I have skepticism about the privacy of it and the evidence based knowledge that it is assuming is accurate when answering a question.* (P69)

4.5 Other Motivations

A small proportion of the participants discussed motivations that did not fit into the four main categories presented above and were not sufficiently frequent to justify separate categories.

Examples of such motivations (all of which were reported by only one participant each) include the following:

- It is easier to talk to a chatbot than to talk to people about important issues.
- Chatbots can provide automatic responses when others are not available.
- One can more easily identify an account as a bot and subsequently block it.
- Chatbots can be a default method of customer support.

Three of the comments coded as "other motivations" were difficult to comprehend.

5 Discussion

We have provided an overview of why people use chatbots and listed participants' responses to an open question regarding their motivations for using chatbots. In this section, we will discuss the findings in terms of U&G. We will then consider the implications of our findings for the future design of chatbots. Finally, we will discuss the study's limitations and possible avenues for future research.

5.1 Productivity is Important

Productivity was the most frequently reported motivation; thus, the majority of chatbot users seek quick and consistent feedback when searching for information or assistance. This finding might reflect the use of chatbots in the customer service domain. This finding may also reflect a general trend for users to gravitate toward immediate communication channels. The broad adoption of private messaging platforms such as Facebook Messenger, WhatsApp, and Snapchat reflects users' interest in more instrumental or goal-directed communication with fewer interruptions compared to regular communication on Facebook and Twitter.

Information has been recognized as an important category of gratification in previous U&G studies, [22]. Yet, the typical chatbot user's need for information may require more immediacy and interactivity than the information needs associated with other media. This hypothesis is in line with recent research identifying young social media users' need for instant gratification. For example, Brandtzaeg et al. [25] suggest that youths communicating with organizations through social media crave immediate feedback and dialogue and action-oriented engagement in order to achieve a clear goal.

Other studies have highlighted the fact that many people, particularly those from Western cultures, seek to spend time productively and may feel guilty when they waste time [26]. Similarly, users in this study often referred to the quick response and productivity of chatbots as key motivations for using them. The need for productivity might be specific to certain cultures, and so it may be worth investigating in this user group.

Instant need for informational feedback may also be related to the concept of usefulness. Usefulness concerns the extent to which a service is perceived as beneficial by performing a specific task quickly and reliably [27]. For chatbots to be successful in the studied user group, they must help users resolve a task or achieve a concrete goal in an effective and efficient manner; in other words, they need to be easy, fast, and convenient. Also, they need to fulfill a valued productivity goal, such as getting help or access to information on the fly.

5.2 Entertainment and Social Motivations Motivate Fewer People but are Important to Some

Entertainment and fun are important aspects of social relations between humans. Likewise, entertainment and socialization may be seen as aspects of the relationship between humans and chatbots. The need for entertainment and a sense of social relationship is also highlighted in U&G and recent U&G studies on social media in particular [24] and online media in general [22].

Many activities in our daily life involve socialization and entertainment. Consequently, Thackara [28] argues that systems should provide users with a social platform or sense of community to generate good user experiences. Similarly, Monk suggests that interactive systems should be designed to support enjoyable social interactions [29]. Sensitivity to this need for entertainment and social relations might be even more important in the context of chatbot design because chatbots are more humanlike than other interactive systems. Thus, users may expect chatbots to be entertaining or social.

Yet, it should be noted that entertainment and social motivations do not exclude productivity motivations. On the contrary, more than one-third of the participants reporting entertainment or social motivations also reported productivity motivations. People want to get their jobs done, but many prefer to do so in a social and enjoyable manner.

5.3 Novelty is a Motivator for Some

Curiosity as a motivation related to news consumption or information-seeking behavior has been identified as a key gratification in previous U&G studies. For example, McQuail [30] argues that "satisfying curiosity and general interest" (p. 87) is a key gratification associated with media use. Such U&G studies have, however, focused on gratifications related to information and content in the context of older mass media such as television and newspapers, not on motivations related to the novelty of interactive technologies such as chatbots.

As discussed in the background of this paper, the curiosity and sense of novelty associated with new technologies and features may be relevant at least for early adopters or innovators, and perhaps specific to these groups. The rest of the population will often view trying out novel technologies as a higher risk and therefore require assurance from trusted peers. This is thoroughly discussed in literature on theory of diffusion of innovations [23], but not in the literature on U&G.

According to Rogers, early adopters and innovators are risk-takers because trying new things may result in failure [23]. Many are interested in novel technologies because of personal entertainment, but early adopters and innovators are more interested in new experiences and learning things before others. For chatbot users, the perceived novelty of chatbots may drive some to use and experiment them. However, to establish a sustainable pattern of usage, chatbots must increase productivity for the late majority for this group to adopt chatbots as a preferred means of interaction.

5.4 Implications for Chatbot Design

A main challenge of user research on this topic is the rapid change in technological developments and user preferences related to chatbots [24]. However, some of the main motivations to use chatbots may be stable over time because they reflect basic needs, such as productivity and social interaction.

Our main findings relate to the key gratifications identified in earlier U&G literature: productivity, entertainment, social and relational purposes, and novelty. The importance of productivity as a motivation for chatbot use is striking, particularly because chatbots for socializing and small talk, such as A.L.I.C.E, Cleverbot, and Mitsuku, have been available for longer than the productivity-oriented chatbots on messaging platforms. Hence, chatbot designers should focus on designing and developing chatbots that are perceived as useful because they provide necessary help or information in an effective and efficient manner. To do so, chatbot designers must identify cases in which chatbots fulfill users' need for productivity more efficiently than what is possible through other methods of interaction. The success of chatbots as personal assistants and health advisors exemplifies the need to design for productivity.

In addition to considering productivity, the chatbot interaction experience can be significantly strengthened by catering to entertainment and social or relational motivations. For example, a productivity-oriented chatbot may benefit from a friendly or empathic appearance. Leading chatbot platforms like Google's Api.ai, include components to support small talk, such as to start a conversation. While the overall purpose of a chatbot may be productivity-oriented, including socialization or entertainment as a feature will be appreciated by a substantial proportion of chatbot users.

The need to balance productivity with entertainment and relational aspects indicates that the relationship between humans and chatbots may be different than the relationships between humans and other tools, such as dishwashers or refrigerators. Thus, chatbots may need to be designed as a tool, toy, and friend.

5.5 Limitations and Future Work

The present study is subject to limitations. First, the chatbot users that participated in this study were both self-selected and filtered by some initial questions. They are, therefore, not representative of the population at large. The participants consisted only of chatbot users, and should therefore be regarded as early adopters, comprising about 14 % of population which, is first to try new ideas, technologies, and services. This may explain why a lot of respondents perceived chatbots to be helpful and efficient. Early adopters may be fundamentally optimistic about future technologies. Hence, this part of the population may focus on the future potential chatbots rather on the current limitations. This also highlight the importance to include a broader part of the population in future studies. However, a strength of the present study is that the sample was large and included users from all over the US. Future studies may benefit from including chatbot users from other countries to determine how users' motivations change across cultures.

Second, the present study involves only a preliminary analysis of the presented data set. We plan to expand the results of this study with additional data collection and analysis in future work. We also plan to investigate how different motivational patterns are linked to age and gender and different chat platforms, as well as analysis of specific chatbots being used. Further, future work should analyze other aspects related to motivations and end-user loyalty, such as why people reduce or stop their use of chatbots.

6 Conclusions

Chatbots potentially represent a new paradigm in how people will interact with data and services in the future. Currently, there is a lack of empirical investigations into why people use chatbots. This study provides needed insight into the motivational factors related to use of conversational interfaces. Its results can guide future research on this topic, which may provide new insights and guide future design and development of chatbots.

Acknowledgment. This study is funded by the research project Human-Chatbot Interaction Design, supported by the Research Council of Norway, IKTPLUSS (p.nr 270940).

References

1. Dale, R.: The return of the chatbots. Nat. Lang. Eng. **22**(5), 811–817 (2016)
2. Følstad, A., Brandtzaeg, P.B.: Chatbots – the new world of HCI. ACM Interactions (2017, in Press)
3. USA Today: Microsoft CEO Nadella: "Bots are the new apps" (2016). https://www.usatoday.com/story/tech/news/2016/03/30/microsof-ceo-nadella-bots-new-apps/82431672/
4. Xu, A., Liu, Z., Guo, Y., Sinha, V., Akkiraju, R.: A new chatbot for customer service on social media. In: Proceedings of the ACM Conference on Human Factors in Computing Systems (2017)
5. Ferrara, E., Varol, O., Davis, C., Menczer, F., Flammini, A.: The Rise of Social Bots. arXiv preprint arXiv:1407.5225 (2014)
6. Simonite, T.: Facebook's Perfect, Impossible Chatbot. MIT Technology Review (2017) https://www.technologyreview.com/s/604117/facebooks-perfect-impossible-chatbot/
7. Coniam, D.: The linguistic accuracy of chatbots: usability from an ESL perspective. Text Talk **34**(5), 545–567 (2014)
8. Malhotra, Y., Galletta, D.F., Kirsch, L.J.: How endogenous motivations influence user intentions: beyond the dichotomy of extrinsic and intrinsic user motivations. J. Manag. Inform. Syst. **25**(1), 267–300 (2008)
9. Weizenbaum, J.: ELIZA—a computer program for the study of natural language communication between man and machine. Commun. ACM **9**(1), 36–45 (1966)
10. Shah, H., Warwick, K., Vallverdú, J., Wu, D.: Can machines talk? Comparison of ELIZA with modern dialogue systems. Comput. Hum. Behav. **58**, 278–295 (2016)
11. Vinyals, O., Le, Q.: A Neural Conversational Model. arXiv preprint arXiv:1506.05869 (2015)
12. Venturebeat (2016). https://venturebeat.com/2016/08/25/niki-ais-new-messenger-bot-lets-you-hail-cabs-and-order-snacks-in-india/
13. James, G.: A Complete Guide to Chatbots (2016). http://www.garethjames.net/complete-guide-chatbots/
14. Hill, J., Ford, W.R., Farreras, I.G.: Real conversations with artificial intelligence: a comparison between human-human online conversations and human-chatbot conversations. Comput. Hum. Behav. **49**, 245–250 (2015)
15. Corti, K., Gillespie, A.: Co-constructing intersubjectivity with artificial conversational agents: people are more likely to initiate repairs of misunderstandings with agents represented as human. Comput. Hum. Behav. **58**, 431–442 (2016)
16. Holtgraves, T.M., Ross, S.J., Weywadt, C.R., Han, T.L.: Perceiving artificial social agents. Comput. Hum. Behav. **23**(5), 2163–2174 (2007)
17. De Angeli, A., Johnson, G.I., Coventry, L.: The unfriendly user: exploring social reactions to chatterbots. In: Proceedings of the International Conference on Affective Human Factors Design, London, pp. 467–474 (2001)
18. Rubin, A.M.: Uses and gratifications. In: Nabi, R.L., Oliver, M.B. (eds.) The SAGE Handbook of Media Processes and Effects, pp. 147–159. Sage, Washington, D.C. (2009)
19. Stafford, T.F., Stafford, M.R., Schkade, L.L.: Determining uses and gratifications for the internet. Decis. Sci. **35**(2), 259–288 (2004)

20. Katz, E., Blumler, J.G., Gurevitch, M.: Utilization of mass communication by the individual. In: Blumler, J.G., Katz, E. (eds.) The Uses of Mass Communications: Current Perspectives on Gratifications Research, pp. 19–32. Sage, Beverly Hills (1974)
21. Joinson, A.N.: Looking at, looking up or keeping up with people?: Motives and use of Facebook. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, pp. 1027–1036. ACM Press (2008)
22. Sundar, S.S., Limperos, A.M.: Uses and grats 2.0: new gratifications for new media. J. Broadcast. Electron. **57**(4), 504–525 (2013)
23. Rogers, E.M.: Diffusion of Innovations. Simon and Schuster, New York (2003)
24. Brandtzæg, P.B., Heim, J.: Why people use social networking sites. In: Ozok, A.A., Zaphiris, P. (eds.) OCSC 2009. LNCS, vol. 5621, pp. 143–152. Springer, Heidelberg (2009). doi:10.1007/978-3-642-02774-1_16
25. Brandtzaeg, P.B., Haugestveit, I.M., Lüders, M., Følstad, A.: How should organizations adapt to youth civic engagement in social media? A lead user approach. Interact. Comput. **28**(5), 664–679 (2016)
26. Foley, C.: The art of wasting time: sociability, friendship, community and holidays. Leisure Stud. **36**(1), 1–20 (2017)
27. Tsakonas, G., Papatheodorou, C.: Exploring usefulness and usability in the evaluation of open access digital libraries. Inf. Process. Manag. **44**(3), 1234–1250 (2008)
28. Thackara, J.: The Design Challenge of Pervasive Computing. CHI (2000). http://www.doorsofperception.com/projects/chi/
29. Monk, A.F.: User-centred design: the home use challenge. In: Sloane, A., van Rijn, F. (eds.) Home Informatics and Telematics: Information Technology and Society, pp. 181–190. Kluwer Academic Publishers, Boston (2000)
30. McQuail, D.: Mass Communication Theory: An Introduction, 2nd edn. Sage, London (1987)

Cascades on Online Social Networks: A Chronological Account

Nora Alrajebah(✉), Thanassis Tiropanis, and Leslie Carr

Web and Internet Science, University of Southampton, Southampton, UK
{N.Alrajebah,tt2,lac}@soton.ac.uk

Abstract. Online social network platforms have served as a substantial venue for research, offering a plethora of data that can be analysed to cultivate insights about the way humans behave and interact within the virtual borders of these platforms. In addition to generating content, these platforms provide the means to spread content via built-in functionalities. The traces of the spreading content and the individuals' incentives behind such behaviour are all parts of a phenomenon known as information diffusion. This phenomenon has been extensively studied in the literature from different perspectives, one of which is cascades: the traces of the spreading content. These traces form structures that link users to each other, where these links represent the direction of information flow between the users. In fact, cascades have served as an artefact to study the information diffusion processes on online social networks. In this paper, we present a survey of cascades; we consider their definitions and significance. We then look into their topology and what information is used to construct them and how the type of content and the platform can consequently affect cascades' networks. Additionally, we present a survey of the structural and temporal features of cascades; we categorise them, define them and explain their significance, as these features serve as quantifiers to understand and overcome the complex nature of cascades.

Keywords: Social network analysis · Information diffusion · Cascades

1 Introduction

Since its emergence, the Internet has created a venue for human-to-human social interaction. In fact, the demand for some form of social networking was raised early on. This was facilitated by different types of computer-mediated communication (CMC), where 'humans' communicated with each other via the "instrumentality of computers" [34]. This comes in many forms on the Internet such as instant messaging, emails, and chat rooms. CMC was the focus of much research in the effects of such communication on social systems. In fact, Kaplan and Haenlein [36] stated that "... the Internet started out as nothing more than a giant Bulletin Board System (BBS) that allowed users to exchange software,

I. Kompatsiaris et al. (Eds.): INSCI 2017, LNCS 10673, pp. 393–411, 2017.
https://doi.org/10.1007/978-3-319-70284-1_31

data, messages, and news with each other." This statement emphasises the core purpose of the Internet as a medium that facilitates different forms of social interactions.

The invention of the Web in 1989 added another dimension to communication on the Internet, providing a wide range of possibilities for human interaction [1,11]. The advance of the Web 2.0 offered a variety of applications that fundamentally changed the way users communicate such as wikis, blogs, RSS, podcasting, and social networks [39]. Therefore, in addition to communication and collaboration, individuals began to contribute to the Web by adding user-generated content. That is what differentiates Web 2.0 from the previous Web [49]. Online social networks have seen a popularity surge following the proliferation of Web 2.0 applications [33]. However, their core purpose is not new; they merely emphasise the Internet's main purpose: facilitating the exchange of information between its users [36]. The Web offers an enormous amount of data that can be analysed to cultivate insights about the way humans behave and interact with each other online. Simultaneously, with the advancements of the Web technologies, a phenomenon has been observed; in the early days on blogosphere, bloggers would share the same URL after being exposed to it by other bloggers, in a cascading manner that can be traced. This phenomenon is information diffusion, and it is concerned with studying the way information is spread on the Web. Online social networks have proven, in many occasions, their vitality for a range of activities that are powered by information diffusion, such as: mass convergence and emergency events [35], spreading information about good practices such as saving energy on earth day [20], bringing people's attention to incidents that might lead to 'public shaming' behaviour [45].

There are three components in any information diffusion process: the content, the context that facilitate the diffusion and the outcome of the process which is the cascade [5]. The earliest research in this field studied diffusion in the blogosphere [4,41]; as new platforms have emerged, they have been used to analyse information diffusion dynamics [3,10,24,38,40]. Research in the field of information diffusion varied according to the purpose of study and the diffusion component(s) that is been taken into consideration. Hence, in their survey of information diffusion in online social networks; Guille et al. [32] categorised the research challenges and approaches in the field into three categories: (1) detecting popular content, (2) modelling information diffusion, and (3) Identifying Influentials.

In this paper, our focus is on cascades, which are defined as the structural representation of the information diffusion and are often perceived as the final outcome of the process [25,42]. Cascades are amplified on online social networks platforms by built-in mechanisms that allow users to share content while crediting the source or the person who posted it [15]. The aim of this paper is to provide a holistic overview of cascades based on their relatively long research history; exploring how cascade-related research has evolved throughout the years. We will attempt to answer the following questions about cascades: *(a) What are they? (b) What is their significance?, (c) How cascade networks are constructed?*

(d) What are the effects of the content type and the available data on cascade construction approaches? (e) What features we can use to analyse cascades? i.e., how can we quantify cascades?

The first two questions set the scene; they emphasise the importance of cascades analysis as a proxy to unveil the sharing dynamics between users on the Web. The answers of questions c, d and e will be presented as a review of two aspects related to cascades: cascade networks construction and cascade features. We will discuss how these two aspects change depending on the content type in the diffusion event and the platform's functionalities. In addition, we will look into the impact of the data that is made available for collection on cascade's analysis. By tracing cascade-related research across platforms, we aim to provide an overview of cascades, which will help designing research problems, and will help researchers as guidelines to construct and analyse cascades.

This paper is organised as follows, Sect. 2 provides some background including cascades definitions, significance and purposes. In Sect. 3, we will look into cascade construction and the different construction approaches. Section 4 is dedicated to cascade features, it is divided into two parts: the structural features and the temporal features. And finally Sect. 5 concludes with some remarks about cascades and how this paper can help researchers who would like to study them.

2 Background

2.1 What are Cascades Networks?

Networks, in their general sense, are structures that consist of a set of nodes and links; the links associate nodes with each other, encapsulating a specific type of a relationship between the two. In mathematical terms, networks are modelled as graphs with vertices and edges [48]. The core concept of networks is their 'connectedness', a phenomenon that has been observed in fields such as Biology, Computer Science and Sociology, and it arises from the flexibility of the definition [22]. A social network can be defined as a network where the nodes represent people and the links represent the relationships and interactions between them [37,48]. Examples of such relationships are: acquaintance, friendship, co-authors, co-workers, affiliation, family relationships, information exchange, etc. [29]. All of these networks link people and, via these links, people interact with each other for many purposes such as: talking to each other, sharing information, and collaborating. One example of such networks are cascade networks, which are networks that link people based on the direction of the flow of information/content between them.

For economists, information cascade occurs when an individual decides that it is optimal to follow the behaviour of those before him after observing their behaviour, without taking into account his own information [13]. In fact, economists differentiate between information cascade and herding behaviour. The difference between the two is that in information cascade individuals decide by making inferences ignoring their own information, while in herding individuals follow the 'herd' without necessarily ignoring their own information [16]. Nevertheless, the term 'cascade' was picked by researchers to describe a similar phenomenon that has been observed in OSNs. In cascades, messages travel through the social network links from one user to another [38]. When gathered, the paths that these messages travel through create a network that resides as a layer on top of the social network. These networks are the cascade networks and the paths messages take are often called information paths [28].

A cascade as defined by Goel et al. [25], comprises a seed individual who shares an item of information independently from any other individual, followed by other individuals who are influenced by the seed to share the same information. Another definition by Leskovec et al. [42] state that cascades are phenomena caused by individuals' influence in which an action or idea becomes widely adopted by others [26,30]; hence, they are known as 'fads' [13]. Both definitions emphasis one point: cascade networks are structures that represent (and preserve) the relationships between users as they share the same content.

2.2 Significance of Cascades

Analysing cascades is an essential step towards understanding the way information propagated on the Internet [21]. When content spreads, it provides us with valuable information about the users involved in the process. As we mentioned in the previous section, cascades represent some form of a relationship between the users. This relationship has been identified in the literature as influence [21]. Identifying influencers has received a significant attention in previous work, and cascades were considered as indicators of influence. Hence, the paths that information takes to reach individuals are recognised as influence paths in many studies, as they directly indicate that one user influenced another to spread the message. In addition to influence, researchers identified another reason behind sharing the same content they do so because there is some degree of homophily among them. Dow et al. [21] state that a user's repeated exposure to a particular content increases the chances of sharing it. They argued that in such a case, these users are subject to both influence and homophily [10]; repeated exposure increases the influence factor, and being surrounded by a group of users who are susceptible to an item means that the user himself is susceptible too. However, cascades do not occur because of influence and homophily only, as both are tied to the nature of content as well. Hence, a cascade informs us about the value of the content itself. Given that users have limited attention, a successful cascade is the one that gets the most attention across the competing cascades at a particular moment [47,53].

Bild et al. [14] refer to cascade networks as implicit networks because they are constructed using a subset of the social network, which they define as an explicit network. They argue that analysing cascade networks is important as these 'implicit' networks can serve as an accurate indication of interest or trust relationships. They conjecture that cascade networks model real-world social, interest and trust networks better than the social network. They argue that connections on the social network (follow/friend) entail that users are willing to listen to each other, but connections on the cascade network are better indicators because they are created using a forceful sharing action that pushes the content to the user's list of friends.

Furthermore, analysing cascades can help detect network evolution and link creation, since users often create new links (follow/friend new users) after being exposed to novel information sources. Myers and Leskovec [46] studied the relation between cascades and the creation of new links in the social network. They related the sudden bursts of connectivity to the dynamics of sharing on Twitter. Antoniades and Dovrolis [8] used the number of retweets and follow reciprocity to model link formation. They also studied link removal dynamics on Twitter after reading a tweet or receiving a retweet from the user. Also, Farajtabar et al. [23] introduced a model that takes into account both activities (sharing and link creation).

To summarise, cascades play an important role in different social network research endeavours:

1. They allow researchers to estimate influence and homophily between users.
2. They work as a proxy to estimate the value of the content that spreads, as successful cascades means that the content attracted a larger number of users.
3. They are better indicators of users' interest and trust relationships than the social network.
4. They help explaining social network evolution and link creation and removal.

2.3 Purposes for Studying Cascades

Cascade studies' purposes vary depending on the objective of the study and the data available for the researchers. Throughout the years, and the different platforms that have been investigated, research purposes have ranged from merely observing and quantifying cascades, to tracking them, predicting information flows, and modelling them [31].

Figure 1 illustrates the four general perspectives for studying cascades. The first, and the essential purpose, is tracking existing cascades then either constructing or inferring them. The ability to construct a cascade depends solely on the data available during the construction process. We will look into this in details in the next section.

The second perspective focuses on quantifying cascades, structurally [21], temporally [31], or just numerically, combined with some platform dependent measures [10]. Often, cascades' tracking is the initial step before quantifying them. For instance, the structural analysis of cascades requires constructing cascades first before the analysis phase can take place. However, some studies focuses on analysing cascades quantitatively, thus, they do not attempt to construct cascade networks as the structure of cascades is not essential for this purpose, e.g. [10]. In Sect. 4 we will present a survey of the structural and temporal features of cascades. We will highlight their significance to understand cascades.

The third perspective looks at modelling cascades, i.e. using generative algorithms to create cascade networks using the characteristics observed from the tracked cascade networks [31,43]. The fourth perspective investigates predictions such as the likelihood that a piece of information will be shared in the first place [50], or the possibility that a popular piece of content will continue to be popular [44], or predicting the future growth of a cascade [18]. Most of the time, one study incorporates one or more purposes in its analysis. However, this paper is focused on the first and second purposes, namely: tracking and quantifying cascades. The third and fourth perspectives are beyond the scope of this paper, thus, we briefly mentioned them here to provide an overall view of the purposes of studying cascades.

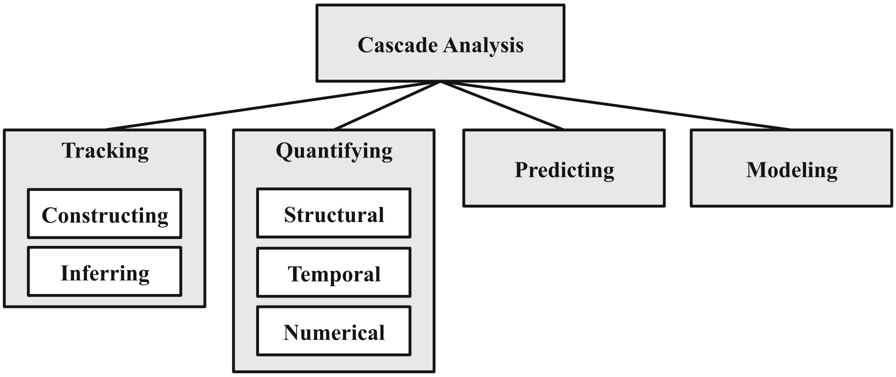

Fig. 1. Purposes for studying cascades

3 Constructing A 'Cascade' Network

As we mentioned earlier, within social networks, many sub-networks can be created using the same nodes that can be linked using edges with various meanings. As soon as information starts to spread within a population, another layer could be added on top of the original network to represent the flow of information [28]. This is often called a diffusion/propagation network or a cascade network [22]. Using Twitter as an example, instead of creating a network of followers, we could create a network where each node represents a user and each link represents a retweet direction. Thus, if A retweeted a tweet posted by B, then there would be a link from B to A, creating what is known as a 'retweet network' [56], or as we will refer to it here a 'cascade network'.

As we mentioned in the previous section, to track existing cascades, they must be either constructed or inferred. In the early studies of cascades, in blogs, for instance, there were no built-in mechanisms for diffusion; thus, most of the early studies used various features to infer cascade networks. Adar and Adamic [4] added a link between two blogs if there is an explicit link to the other. If there is no explicit link, they infer it using a number of features related to the blog network structure, historical data about the blogs' posts, text similarity, and timestamps. Most early studies of cascades on online social networks exploited users' written credit attribution of content sources to infer cascade networks. Examples of credit attributions are "RT", "via", "retweet", and "reshare" [21]. There were also many attempts to use the social network and timestamps to infer cascade networks [27]. However, with more contextual information available it is possible to construct more accurate cascade networks. For instance, Dow et al. [21] used information about reshares, timestamps and clicks on feed, to infer cascade networks and compare them with cascade networks constructed solely from tracked information.

More recently, online social network platforms start incorporating the ability to share content with a click on a button; for example retweet on Twitter, Reblog on Tumblr and Share on Facebook. With these functionalities in place, users can share different types of content easily. As a consequence, tracking existing cascades is now feasible with the appropriate access to data. Thus, researchers are now able to construct existing cascades directly from the platform.

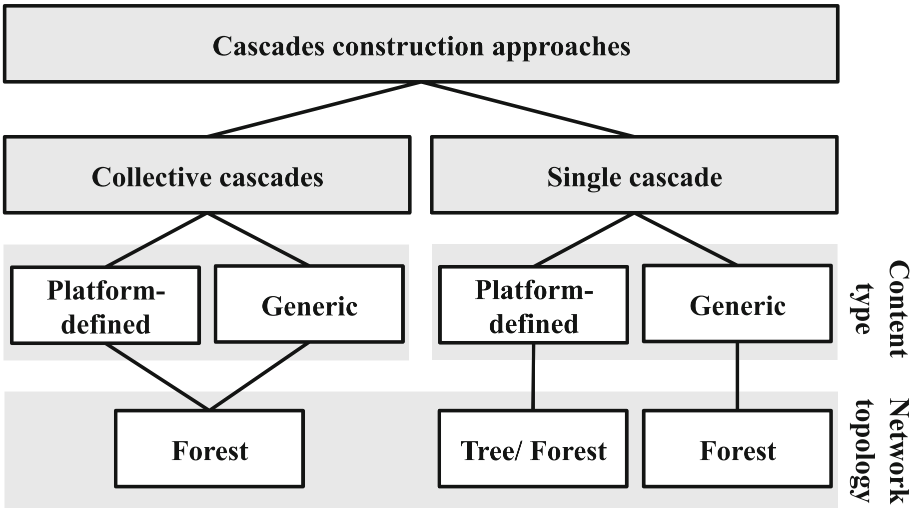

Fig. 2. Cascade networks construction approaches and their resulted topologies

3.1 Cascade Networks Topologies

A cascade is often perceived as a tree that has a single root (the cascade initiator) which is linked to other nodes. Further nodes can be added by linking to the existing nodes in the cascade network and all of the added links follow a strict time order [41]. However, cascades are not always shaped as trees, in fact, their structure changes depending on the type of content these cascades networks represent. Anderson et al. [7] classified cascade networks into: information-sharing networks in which information spread between the users and signups which mimic the adoption of a new technology. This classification does not specify the topology of the generated cascade network. Thus, here we present a different classification of cascade networks based on their topology. The basis of this classification is the content type and the diffusion mechanism provided by the platform.

There are two main approaches to constructing cascade networks that have been used in research. Figure 2 illustrates them and the resulting cascades' topologies generated by each approach. The first approach is collective cascades, in which a large cascade network is constructed, linking users according to their sharing activities (retweet/reblog) collectively for a group of cascading items. The topology of this network is a forest that has several components. These large networks are useful to study the sharing activity patterns within a platform [14,54]. Collective cascade networks are often weighted to represent how often a link occurs between two nodes [41].

The second approach is for single cascades in which cascade networks are constructed for each item that has been shared separately. Of the two categories of content, the first is a platform-defined elements such as a tweet in Twitter or a post in Tumblr. The second category (generic elements) covers any element that can be embedded within platform-defined elements such as a URL, a hashtag, a text, or a photo. Different content types require different data collection and analysis methods, and they create a completely different network topology.

The platform-defined elements that can be shared are for example: a post on Tumblr and Facebook, or a tweet on Twitter. This type of content spreads via explicit diffusion functionalities such as retweeting, sharing or reblogging. Their spread generates cascades that can be tracked or inferred on the platform. Cascades are constructed from the flow of information from users who might or might not be connected to each other by a relationship within the social graph [5,6,18,19]. These cascade networks ideally follow a tree topology; the root is the source (author) and from there content travels across the social network. However, in many cases due to the limited access to the platform, some data might be missing because it is deleted, the topology of the generated cascade network will be a forest where there will be separate components for each isolated part that can not be linked to the main tree due to missing data [6,52].

Because the diffusion of generic elements, such as hashtags and URLs, does not occur via explicit diffusion functionalities in social networks. Thus, timestamps are often utilised as an indicators of diffusion between users assuming that these users have an established social relationship in the social network graph. Cascade networks of generic items are different to cascade networks of one story. These networks incorporate multiple introductions of the same item in the network, thus naturally their topology will be a forest with separate components (sub-cascades). Hence, the number of sub-cascades and their sizes can be used as structural features of these networks [24].

Collective cascades networks can be easily converted into single cascade networks by separating the different branches of the network where they are related to the same story (message). For instance, Leskovec et al. [41], generated cascade networks following the two approaches from blogosphere. They constructed a post network that links posts if they credit each other. From the post network they constructed a blog network by collapsing the links between blogs and assigning weights to them. Following this method, they constructed separate cascade trees from the post network. Sections 3.3 and 3.4 will discuss the cascade construction approaches used in different platforms, including the data used for their construction, the detected diffusion mechanism, and the structure of the cascade networks.

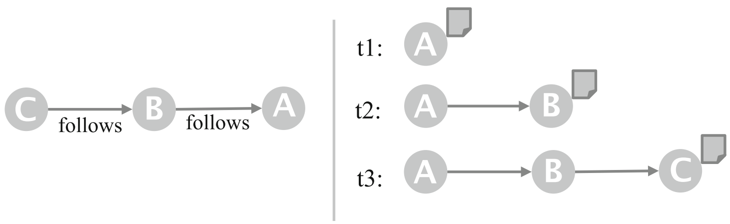

Fig. 3. Link types in cascade networks, left: relationship perspective, right: information flow perspective

3.2 Link Direction

Edges between the nodes in a network might convey different meanings. For instance, Bild et al. [14] identify the number of users who retweet from a user as the popularity; while the prolificity refers to the number of users a user retweet from. Hence, the direction of edges in a network can have different meanings. Consequently, all the measures that rely on the edges' direction will be affected.

Figure 3 illustrates two possible uses for edges' direction as used in the literature For example, suppose that we have three users, A, B and C. For simplicity, suppose that we have the following settings: user B follows user A and user C follows user B. Then, each time user A posted some content user B will be exposed to it and when user B shares that content after seeing it; user C will be exposed to the content too and can share it as well. In such a scenario, there are two possible representations:

1. **Relationship perspective:** If our concern is to represent who is linked to whom i.e., who follows whom, then the in-link from B to A means that B is linked to A, and the in-link from C to B means that C is linked to B. This is shown on the left in Fig. 3, this representation is often referred to as the social network or the follow network.
2. **Information flow perspective:** In this case, the in-link from one user (A) to another (B) means that B is exposed to whatever information A has and when B shares that information too an edge will be drawn from A to B indicating the flow of information from A to B. This representation is often used for cascade networks. Figure 3 shows how this network can be constructed

cumulatively at different timestamps. At timestamp t1, A posted a content, then when B was exposed to it, B decides to share it at timestamp t2, hence the edge from A to B and so on.

3.3 Cascades in Blogs, Recommendation Networks and Internet-Chain Letters

Data used to construct/infer cascades
As mentioned earlier, in the early days of blogosphere there were no convenient mechanism to share content. Thus, instead of following the traces, cascades are inferred using a variety of measures such as: posts text, explicit links to other blogs, features about the blogs network, the blog and the timestamps [4]. In another study of cascades on blogs, the In-links/out-links between blog posts and timestamps were utilised to construct cascade networks [41].

On the other hand, on recommendation networks information about: products, time of recommendation, whether the product is purchased, and time of purchase are utilised to infer these networks [42]. Also, Liben-Nowell and Kleinberg [43] used the ordered list of users who forwarded the petition to construct the cascades of chain-letters.

Diffusion mechanism
As we can see the lack of explicit diffusion functionality means that various mechanisms of diffusion were identified such as: posting a URL in a blog [4], recommending a product [42], linking between posts on blogs [41], and forwarding of a petition letter from one user to another [43].

Cascade networks topology and components
The network topology of these cascades and their components vary based on the platform and the purpose of them. For instance, in [4] the cascade networks structure is trees, where the nodes are blogs and the edges between them are inferred to show the direction of diffusion of information between the blogs. While Leskovec et al. [42], constructed a posts network that links posts in different blogs, and a blogs network which is a collapsed and weighted version of the posts network. Both networks are forests and they extracted separate cascade trees from the posts network. On recommendation networks a separate group networks and a product networks are constructed, where the nodes are the customers and the edges connect customers' product recommendations [41]. Finally, in the work Liben-Nowell and Kleinberg [43] the lists of users in each petition contains duplicates or missing users. Thus, the cascade networks are trees inferred by removing edges that did not appear in a sufficient number of copies. Thus, the nodes are users and the edges represent the direction of information flow between them.

3.4 Cascades in OSNs

Data used to construct/infer cascades
Depending on the content type in each study and the diffusion mechanism, the data needed to construct cascade networks on OSNs vary from: retweets on

Twitter [12,14,38], reblogs on Tumblr [5,6,17,54], share on Facebook [18,19,21]. The tweet texts, timestamps and social network are used in [24] to infer cascade networks of URLs. Yang and Counts [55] analysed tweets' texts that contain topics and mentions of other users to construct cascades. Also, text analysis (status updates that include the meme and the words 'copy', 'paste' and 'repost'), lists of users who commented on users' status and timestamps are used in [3] to construct cascades of memes on Facebook. In another study of cascades of memes on Facebook, the social network, time, text similarity measures are used [2]. On LinkedIn signups and timestamp are used to construct cascade networks of invitations [7]. These examples shows the diverse views of cascades on OSNs; they show us how the diffused content type affects the cascade, and the varieties of data that can be used to either construct or infer cascade networks.

Diffusion mechanism
On OSNs the main diffusion mechanism is provided by a platform's functionality (retweet, reblog, share). Other mechanisms of diffusion are: posting a URL [24], or crediting the source using 'RT @' in tweet text [14,24]. For memes, the diffusion mechanism is simply copy and paste of textual memes [2,3].

Cascade networks topology and components
Various cascade networks topologies are used based on the content type, as mentioned earlier platform-defined elements generate trees, while generic elements generate forests. For example, Kwak et al. [38] created retweet trees for each tweet in their dataset and forests for each topic. Also, in [24], because the diffusion mechanism used is either posting a URL or crediting the source, the generated cascades' structure is a forest. Due to their nature, cascades of memes are forests [2,3]. There are also two studies that constructed large cascade networks of collective cascades [14,54]. In general, the nodes in most of the cascades on OSNs are users, and the edges always indicate the direction of information flow between them. An exception was found in [2], where the nodes are meme variants and the edges between them link a meme variant to its parent.

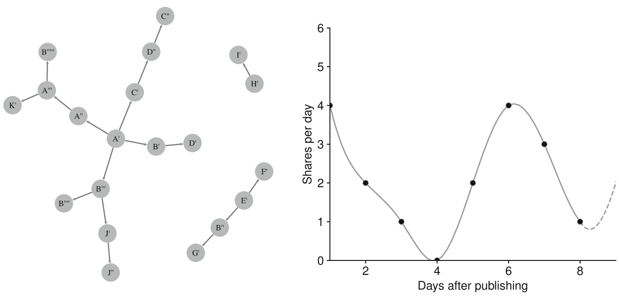

(a) (b)

Fig. 4. Cascade data dimensions, (a) cascade network data, (b) time-series data

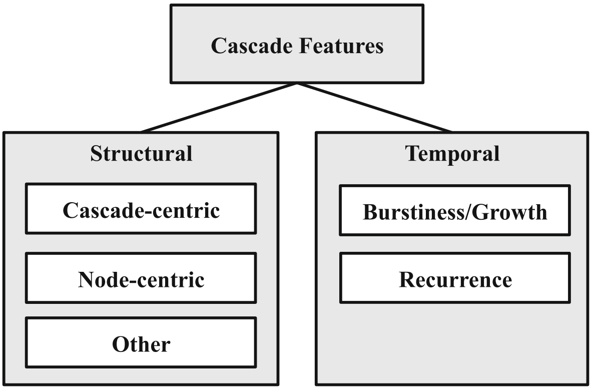

Fig. 5. Cascade features classification

4 Cascades Features

In general, the data we can harvest about cascades is multidimensional in its nature. It has a twofold purpose: the first is to allow cascade networks to be constructed using the detailed information about users sharing from other users; the second is to allow the creation of a time series dataset, where the number of sharing activities at a given time (day or hour) after publishing is recorded. Figure 4 illustrates these two data representations that are used for the analysis. The first is linked to the relation between the users involved in the cascade, i.e. who influenced whom to spread the content. The time-series information about cascades provides the number of diffusion events that occur at a given time. Each of these dimensions is related to a different aspect: either the structural or the temporal. These two aspects complement each other and provide a better understanding of cascades, as Scott [51] argued that the temporal aspect adds value to the structural aspect when analysing data from social networks.

The level of access researchers have to the platform's data determines the type of data they can gather. For instance, utilising a privileged access ensures that both dimensions are harvested, minimising the effect of missing or deleted data. In addition, with privileged access researchers have unlimited access to rich metadata such as clicks in News Feed [21]. As a result, they can infer cascades more accurately. Figure 5 illustrates the two classes of cascades features; in the next subsections, we will explore the structural and temporal features of cascades. We will identify these features and highlight their significance in relation to cascades' analysis.

4.1 Structural Features

Analysing the structural features of cascades includes studying their structure and quantifying cascade networks' properties. According to Liben-Nowell and

Table 1. Structural features

Feature	Definition
Cascade-centric features	
Depth, range, distance to the root, Maximum depth, maximum hop count, cascade height [3,6,9,17,21,24,25,38,43,55]	Represents the height of a cascade, it is calculated using the number of subsequent occurrences of message passing events, i.e. maximum number of hops or range of influence. Maximum depth and average depth can be measured too
Width [43]	The maximum size of a set of nodes which share the same depth
The fraction of nodes with exactly one child [6,43]	How many nodes in a cascade with exactly one child
Scale [6,55]	The number of nodes influenced by a message at depth equals one
Wiener index [5,7,18]	It is used to measure the structural virality of a cascade. It is computed as the average distance between all pairs of nodes in a cascade
The percentage of adoption per depth [6,7,25]	The percentage of adoption events that occur at each depth from the root
Number of nodes at depth = 1 [6,21]	The number of nodes that are one step away from the author
Connectivity rate [52]	The percentage of users who have one edge at least, hence they were influenced by other users
Root fragment rate [52]	The percentage of nodes that have either direct or indirect connection to the root node
Diameter [42,52]	The diameter of a network
Node-centric features	
Fanout, branching factor [5,6,21,31]	The number of subsequent cascades that follow directly from a particular node (user)
Size of Sub-cascade [6,21,24]	The size of the sub-cascade under a particular node
Other	
Frequency of distinct cascade structures [17,25,41,42]	After constructing all of the cascade trees in a dataset, it is possible to compute the frequency of cascade structures. This process is computationally expensive as it aggregates all the generated cascade trees to identify similar structures, e.g. trees with root only, or trees with a root and two child nodes

Kleinberg [43], a better understanding of the properties of the structure of cascades leads to better dissemination models. Table 1 lists the structural features of cascades that are categorised into three categories. The first category are cascade-centric features; these features are computed on the cascade level as a whole. The significance of each of these features is as follow:

1. **Depth, range, distance to the root:** Indicates the shape of a cascade, and how far it travels away from the source within the network. When all distances to the root are gathered, they can help assessing whether a cascade is shallow or deep [3,6,9,17,21,24,25,38,43,55].
2. **Width:** Indicates the extent to which a cascade is narrow or wide. It gives hints about the factors that make a message quite popular at one stage within the cascade [43].
3. **The fraction of nodes with exactly one child:** Indicates missing or unsuccessful cascade event [6,43].
4. **Scale:** Indicates how popular/interesting a message gets soon after its first appearance [6,55].
5. **Wiener index:** Gives an indication of the cascade shape, the higher the Wiener index, the more viral the cascade. Cascades with low Wiener index resemble a star shape, where there are few hubs that create the cascade. The Wiener index increases with the increase in cascade size [5,7,18].
6. **The percentage of adoption per depth:** Counting the percentage of adoptions within one degree of a root could indicate whether epidemic-like cascade occurs in the dataset, i.e. if the majority of adoptions recorded in the dataset are within the first few degrees from a root, then one could conclude that most cascades are shallow and small [6,7,25].
7. **Number of nodes at depth = 1:** Nodes (users) at depth 1 are the ones who share directly from the author, meaning that they were exposed to the author's post directly. It might be that they arrive via external resources or direct links. Although there is a possibility that users click on the original post and share from the author rather than from user they receive the post from [6,21].
8. **Connectivity Rate:** Shows whether an edge exists between any two nodes in the cascade. It is useful to examine whether users get their information from the social links (i.e. explicit links via following) if this information was taken into account while constructing the cascade tree [52].
9. **Root Fragment Rate:** Shows whether each node in the cascade is actually linked to the root or not. It is useful to examine whether users get their information from social links (i.e. explicit links via following) if this information was taken into account while constructing the cascade tree [52].
10. **Diameter:** Shows whether cascades are deep or shallow [42,52].

The second category is node-centric structural features, which are computed on nodes level. There are two features in this category: the branching factors and the subcascade size and they both measure individual's influence on the overall cascade [5,6,21,24,31]. However, there is a difference between the two, as the branching factor estimates the immediate influence, the subcascade size estimates the overall influence of one individual on the cascade. It is important to take the two measures into account as some nodes might have a small branching factor but their subcascade might be very large [6].

The last structural feature is the frequency of distinct cascade structures. It helps to detect if there is a repeated cascade pattern, which can be investigated later. When combined with depth, it could help draw some conclusions about the shape of the cascade and how far it branches [17,25,41,42].

Table 2. Temporal features

Feature	Definition
Time passed since message published [21]	How many times a particular message has been passed in relation to the time since it was published
Speed [55]	Detecting whether and when the first cascade will occur (depth = 1)
Time lag between posting and first reshare, elapsed time [6,17,38]	The difference between posting time and the first reshare
Time lag between two sharing events [38]	The difference between two nodes in a cascade
The number of spikes/peaks [19,31]	Spikes refer to high-volume of cascading activities that occur in a short period during the lifetime of a cascade
Cascading density throughout lifetime [6,31]	The timeline of a cascade, it shows the number of cascading activities per day
Maximum time between reshares [6,19]	The maximum time difference between reshares
Cascade growth/cascade popularity [3,6,7,21,41]	The relation between the growth in cascade size through time. The rate at which cascades gain their size (i.e. popularity)
Recurrence [19]	Recurrence occurs if a cascade has at least two peaks in addition to other conditions

4.2 Temporal Features

There are two approaches to analyse the temporal aspect of cascades. The first tracks and describes existing cascades' temporal features, e.g. how fast information spreads, for how long trendy content keeps its popularity, and the overall growth of cascades over time, such as: whether cascades show patterns like 'burstiness' or sparks. The other line of research uses cascade's temporal patterns to either predict or model the cascade's future popularity. Most of these studies do use the word 'cascade', because they are concerned with the temporal aspect of the diffusion of online content. However, the underlying structure of online content diffusion is an implicit cascade network.

Table 2 lists a number of cascades' temporal features, their significance is as follow:

1. **Time passed since message published:** Shows the growth of cascade and the fade of interest in the message over time [21].
2. **Speed:** Indicates how fast users would be influenced to spread the message or generally react using other means of interaction like reply or mention [55].
3. **Time lag between posting and first reshare, elapsed time:** Measures the resharability of content: the larger the lag the less likely a content will be reshared [6,17,38].
4. **Time lag between two sharing events:** Shows the speed at which a cascade occurs in relation to the distance between nodes, i.e. sharing events [38].
5. **The number of spikes/peaks:** Measures the degree to which a cascade provokes high volume of cascading during its lifetime [19,31].
6. **Cascading density throughout lifetime:** Helps assessing the temporal patterns of diffusion, whether it has spikes or maintains a steady growth. [6,31]
7. **Maximum time between reshares:** Indicates the maximum idleness period within a cascade [6,19]
8. **Cascade growth/cascade popularity:** Helps to show whether a cascade size grows linearly as time passes or in different ways. This helps detect whether the growth in cascade size occurs in short intervals or whether it grows with time. It also shows the periods of idleness and spikes in the cascade timeline [3,6,7,21,41].
9. **Recurrence:** Helps identifying cascades that regain their popularity after a period of idleness [19].

5 Conclusions

In this paper, we presented a survey of cascades, cascade networks and cascade features. Our aim was to investigate these subjects while considering two aspects: the content type and the platform. The main message this paper conveys is that content type is significant in the process of constructing and analysing cascades. Not only that, but, content type has an impact on the approaches used to collect the datasets as well. In addition, the survey of cascade features will be useful for researchers who would like to study cascades, as it will give them an overall overview of the measures that have been used in the literature including their significance as cascades estimators. These features can be used for several purposes related to cascades: quantifying, modelling and predicting their future growth.

References

1. World Wide Web Timeline (2014). http://www.pewinternet.org/2014/03/11/world-wide-web-timeline

2. Adamic, L., Lento, T., Adar, E., Ng, P.: Information evolution in social networks. In: Proceedings of the Ninth ACM International Conference on Web Search and Data Mining (WDSM 2016), pp. 473–482. ACM, New York (2016)
3. Adamic, L.A., Lento, T.M., Fiore, A.T.: How you met me. In: Proceedings of the 6th International AAAI Conference on Weblogs and Social Media (ICWSM) (2012)
4. Adar, E., Adamic, L.: Tracking information epidemics in blogspace. In: Proceedings of the 2005 IEEE/WIC/ACM International Conference on Web Intelligence (WI 2005), pp. 207–214. IEEE Computer Society (2005)
5. Alrajebah, N.: Investigating the structural characteristics of cascades on Tumblr. In: Proceedings of the 2015 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, (ASONAM 2015), pp. 910–917. ACM (2015)
6. Alrajebah, N., Carr, L., Luczak-roesch, M., Tiropanis, T.: Deconstructing diffusion on Tumblr: structural and temporal aspects. In: Proceedings of the 9th ACM Conference on Web Science. ACM, in press
7. Anderson, A., Huttenlocher, D., Kleinberg, J., Leskovec, J., Tiwari, M.: Global diffusion via cascading invitations: structure, growth, and homophily. In: Proceedings of the 24th International Conference on World Wide Web (WWW 2015), pp. 66–76. ACM (2015)
8. Antoniades, D., Dovrolis, C.: Co-evolutionary dynamics in social networks: a case study of Twitter. Comput. Soc. Netw. **2**, 14 (2015)
9. Bakshy, E., Hofman, J., Mason, W.A., Watts, D.J.: Everyone's an influencer: quantifying influence on Twitter. In: Proceedings of the Fourth ACM International Conference on Web Search and Data Mining (WSDM 2011), pp. 65–74. ACM (2011)
10. Bakshy, E., Rosenn, I., Marlow, C., Adamic, L.: The role of social networks in information diffusion. In: Proceedings of the 21st International Conference on World Wide Web (WWW 2012), pp. 519–528. ACM, Lyon (2012)
11. Berners-Lee, T.: Weaving the Web: The Original Design and Ultimate Destiny of the World Wide Web by its Inventor. HarperInformation (2000)
12. Bhattacharya, D., Ram, S.: Sharing news articles using 140 characters: a diffusion analysis on Twitter. In: Proceedings of the 2012 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, (ASONAM 2012), pp. 966–971. IEEE Computer Society (2012)
13. Bikhchandani, S., Hirshleifer, D., Welch, I.: A theory of fads, fashion, custom, and cultural change as informational cascades. J. Polit. Econ. **5**, 992–1026 (1992)
14. Bild, D.R., Liu, Y., Dick, R.P., Mao, Z.M., Wallach, D.S.: Aggregate characterization of user behavior in Twitter and analysis of the retweet graph. ACM Trans. Internet Technol. **15**(1), 24 (2015)
15. Boyd, D., Golder, S., Lotan, G.: Tweet, Tweet, Retweet: conversational aspects of Retweeting on Twitter. In: 2010 43rd Hawaii International Conference on System Sciences (HICSS), pp. 1–10. IEEE Computer Society (2010)
16. Çelen, B., Kariv, S.: Distinguishing informational cascades from herd behavior in the laboratory. Am. Econ. Rev. **94**(3), 484–498 (2004)
17. Chang, Y., Tang, L., Inagaki, Y., Liu, Y.: What is Tumblr: a statistical overview and comparison. SIGKDD Explor. **16**(1), 21–29 (2014)
18. Cheng, J., Adamic, L.A., Dow, P.A., Kleinberg, J., Leskovec, J.: Can cascades be predicted? In: Proceedings of the 23rd International Conference on World Wide Web (WWW 2014), pp. 925–935. ACM, Seoul (2014)
19. Cheng, J., Adamic, L.A., Kleinberg, J., Leskovec, J.: Do cascades recur? In: Proceedings of the 25th International Conference on World Wide Web (WWW 2016), pp. 671–681. ACM (2016)

20. Cheong, M., Lee, V.: *Twittering* for earth: a study on the impact of microblogging activism on earth hour 2009 in Australia. In: Nguyen, N.T., Le, M.T., Świątek, J. (eds.) ACIIDS 2010. LNCS, vol. 5991, pp. 114–123. Springer, Heidelberg (2010). doi:10.1007/978-3-642-12101-2_13
21. Dow, P., Adamic, L., Friggeri, A.: The anatomy of large Facebook cascades. In: Proceedings of the Seventh International Conference on Weblogs and Social Media, (ICWSM), pp. 145–154. AAAI, Cambridge (2013)
22. Easley, D., Kleinberg, J.: Networks, Crowds, and Markets: Reasoning About a Highly Connected World. Cambridge University Press, Cambridge (2010)
23. Farajtabar, M., Gomez-Rodriguez, M., Wang, Y., Li, S., Zha, H., Song, L.: Co-evolutionary dynamics of information diffusion and network structure. In: Proceedings of the 24th International Conference on World Wide Web (WWW 2015), pp. 619–620. ACM (2015)
24. Galuba, W., Aberer, K.: Outtweeting the Twitterers - predicting information cascades in Microblogs. In: Proceedings of the 3rd Wonference on Online Social Networks (WOSN 2010), pp. 1–9. USENIX Association, Boston (2010)
25. Goel, S., Watts, D., Goldstein, D.: The structure of online diffusion networks. In: Proceedings of the 13th ACM Conference on Electronic Commerce (EC 2012), vol. 1, pp. 623–638. ACM (2012)
26. Goldenberg, J., Libai, B., Muller, E.: Talk of the network: a complex systems look at the underlying process of word-of-mouth. Mark. Lett. **12**(3), 211–223 (2001)
27. Gomez Rodriguez, M., Leskovec, J., Krause, A.: Inferring networks of diffusion and influence. In: Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - (KDD 2010), pp. 1019–1028. ACM (2010)
28. Gomez Rodriguez, M., Leskovec, J., Schölkopf, B.: Structure and dynamics of information pathways in online media. In: Proceedings of the Sixth ACM International Conference on Web Search and Data Mining (WSDM 2013), p. 23. ACM (2013)
29. Grabner-Kräuter, S.: Web 2.0 social networks: the role of trust. J. Bus. Ethics **90**(Suppl. 4), 505–522 (2009)
30. Granovetter, M.S.: Threshold models of collective behavior. Am. J. Sociol. **83**(6), 1420–1443 (1978)
31. Gruhl, D., Guha, R., Liben-Nowell, D., Tomkins, A.: Information diffusion through Blogspace. In: Proceedings of the 13th International Conference on World Wide Web (WWW 2004), pp. 491–501. ACM (2004)
32. Guille, A., Hacid, H., Favre, C., Zighed, D.: Information diffusion in online social networks: a survey. SIGMOD Rec. **42**(2), 17–28 (2013)
33. Heidemann, J., Klier, M., Probst, F.: Online social networks: a survey of a global phenomenon. Comput. Netw. **56**(18), 3866–3878 (2012)
34. Herring, S.C.S.: Computer-Mediated Communication: Linguistic, Social, and Cross-Cultural Perspectives, vol. 39. John Benjamins Publishing, Amsterdam (1996)
35. Hughes, A.L., Palen, L.: Twitter adoption and use in mass convergence and emergency events. Int. J. Emergency Manag. **6**, 248 (2009)
36. Kaplan, A.M., Haenlein, M.: Users of the world, unite! The challenges and opportunities of social media. Bus. Horiz. **53**(1), 59–68 (2010)
37. Kempe, D., Kleinberg, J., Tardos, E.: Maximizing the spread of influence through a social network. In: Proceedings of the Ninth ACM International Conference on Knowledge Discovery and Data Mining (KDD 2003), pp. 137–146. ACM (2003)
38. Kwak, H., Lee, C., Park, H., Moon, S.: What is Twitter, a social network or a news media? In: Proceedings of the 19th International Conference on World Wide Web (WWW 2010), pp. 591–600. ACM (2010)

39. Lai, L.S.L., Turban, E.: Groups formation and operations in the web 2.0 environment and social networks. Group Decis. Negot. **17**(5), 387–402 (2008)
40. Lerman, K., Ghosh, R.: Information contagion: an empirical study of the spread of news on Digg and Twitter social networks. In: Fourth International AAAI Conference on Weblogs and Social Media, pp. 90–97. AAAI (2010)
41. Leskovec, J., McGlohon, M., Faloutsos, C., Glance, N., Hurst, M.: Patterns of cascading behavior in large blog graphs. In: Proceedings of the 2007 SIAM International Conference on Data Mining, pp. 551–556. SIAM (2007)
42. Leskovec, J., Singh, A., Kleinberg, J.: Patterns of influence in a recommendation network. In: Ng, W.-K., Kitsuregawa, M., Li, J., Chang, K. (eds.) PAKDD 2006. LNCS, vol. 3918, pp. 380–389. Springer, Heidelberg (2006). doi:10.1007/11731139_44
43. Liben-Nowell, D., Kleinberg, J.: Tracing information flow on a global scale using Internet chain-letter data. Proc. Natl. Acad. Sci. **105**(12), 4633–4638 (2008)
44. Ma, Z., Sun, A., Cong, G.: On predicting the popularity of newly emerging hashtags in Twitter. J. Am. Soc. Inform. Sci. Technol. **64**(7), 1399–1410 (2013)
45. McBride, K.: Journalism and public shaming: Some guidelines (2015). http://www.poynter.org/2015/journalism-and-public-shaming-some-guidelines/326097/
46. Myers, S., Leskovec, J.: The Bursty dynamics of the Twitter information network. In: Proceedings of the 23rd International Conference on World Wide Web (WWW 2014), pp. 913–923. ACM (2014)
47. Myers, S.A., Leskovec, J.: Clash of the contagions: cooperation and competition in information diffusion. In: Proceeding of IEEE 12th International Conference on Data Mining, pp. 539–548. IEEE (2012)
48. Newman, M.E.J.: Networks: An Introduction. Oxford University Press, Oxford (2010)
49. O'Reilly, T.: What is web 2.0 (2005). http://www.oreilly.com/pub/a/web2/archive/what-is-web-20.html
50. Petrovic, S., Osborne, M., Lavrenko, V.: RT to win! Predicting message propagation in Twitter. In: Proceedings of 5th International Conference on Weblogs and Social Media (ICWSM), pp. 586–589. AAAI (2011)
51. Scott, J.: Network analysis. In: Darity, W.A. (ed.) International Encyclopaedia of the Social Sciences. Macmillan, New York (2008)
52. Taxidou, I., Fischer, P.M.: Online analysis of information diffusion in Twitter. In: Proceedings of the Companion Publication of the 23rd International Conference on World Wide Web Companion (WWW Companion 2014), pp. 1313–1318. ACM (2014)
53. Weng, L., Flammini, A., Vespignani, A., Menczer, F.: Competition among memes in a world with limited attention. Sci. Rep. **2**, 1–9 (2012)
54. Xu, J., Compton, R., Lu, T.C., Allen, D.: Rolling through Tumblr : characterizing behavioral patterns of the Microblogging platform. In: Proceedings of the 2014 ACM Conference on Web Science (WebSci 2014), pp. 13–22. ACM (2014)
55. Yang, J., Counts, S.: Predicting the speed, scale, and range of information diffusion in Twitter. In: Proceedings of 4th International Conference on Weblogs and Social Media (ICWSM 2010), pp. 355–358. AAAI (2010)
56. Yang, L., Sun, T., Zhang, M., Mei, Q.: We know what@ you# tag: does the dual role affect hashtag adoption? In: Proceedings of the 21st International Conference on World Wide Web (WWW 2012), pp. 261–270. ACM (2012)

Semantic Social Networks: A New Approach to Scaling Digital Ethnography

Alberto Cottica[1,2], Amelia Hassoun[3], Jason Vallet[4], and Guy Melançon[4(✉)]

[1] Edgeryders, Tallin, Estonia
alberto@edgeryders.eu
[2] University of Alicante, Alicante, Spain
[3] University of Oxford, Oxford, UK
amelia.hassoun@gmail.com
[4] Université de Bordeaux, CNRS UMR 5800 LaBRI, Bordeaux, France
{jason.vallet,Guy.Melancon}@u-bordeaux.fr

Abstract. We propose a data-based approach to doing ethnographic research in a digital environment. It has three main components. First, it treats online conversational environments as human communities, that ethnographers can engage with as they would in onsite fieldwork. Second, it represents those conversations, and the fieldnotes made by researchers thereon, in network form. We call these networks *semantic social networks*, as they incorporate information on social interaction and their meaning. They encode a map of the associations between key concepts as perceived by informants as a group. Third, it uses methods borrowed from network science to process these data.

We present an application of this method to a large online conversation about community provision of health and social care, and discuss its potential for harnessing collective intelligence.

Keywords: Digital ethnography · Network science · Collective intelligence

1 Introduction

The Internet Age has brought about a wave of exploration and innovation into digital ethnographic research methods. Substantial work has been devoted to methods that mine social networking platforms for user-generated content to analyse, often automatically. We propose an alternative approach based on convening an online conversation on the topic of interest. Such conversations function as virtual communities [17]. As such, they lend themselves to participant observation.

The digital nature of the conversational medium transforms the ethnographic evidence into structured data. This offers two opportunities. First, it allows ethnographers to do quantitative analysis on their own qualitative analysis. Second, as quantitative analysis functions as an aggregation layer, it allows ethnographers to handle larger volumes of evidence in coherent, replicable ways.

I. Kompatsiaris et al. (Eds.): INSCI 2017, LNCS 10673, pp. 412–420, 2017.
https://doi.org/10.1007/978-3-319-70284-1_32

In what follows, we show how ethnographic evidence maps onto a type of network that encodes both social interaction and semantic association. We call these networks *semantic social networks (SSNs)*, and claim that a methodology based on them is highly accountable to ethnography as a discipline. Its steps, save for the final quantitative analysis layer, carry naturally over from onsite field research to the digital domain. So, then, does ethnography's distinctive focus on groups of humans and their worldviews.

Additionally, SSNs are much more scalable than traditional ethnography. Coupled with open data and open standards, they can work well with thousands of informants - a scale large enough for most applications, and much larger than that achieved by traditional ethnography. In sum, SSNs have the potential to evolve in a research method that (a) discovers *collective* worldviews of groups of humans; (b) can address open questions (unlike surveys) and (c) scales reasonably well (unlike onsite ethnography). They show promise as tools to harness *collective intelligence*, the processing of information by connected groups of humans [14].

In what follows, we describe the approach and present the results obtained by applying it to a digital ethnography dataset. We first introduce a data model for SSNs (Sect. 2). Next, we present data in SSN form from a study on community-provided health and social care services (Sect. 3). We then illustrate how we used SSNs to aggregate and navigate a large corpus of ethnographic evidence (Sect. 4). Finally, we reflect on some possible extensions to SSNs and their potential for allowing digital ethnography to scale, while still maintaining its methodological advantages (Sect. 5).

2 Semantic Social Networks: A Generalised Data Model for Digital Ethnography:

Ethnography is a qualitative research technique aimed at discovering how a certain group of humans perceives a set of issues. Its unique value lies in that its findings encode the culture and worldview of the group being studied. This makes it especially suited for applications like foresight [1] and democratic stakeholder dialogue [7,21], where social and cultural meanings that arise organically from human interactions are the main objects of research rather than pre-conceived, researcher-imposed analytical categories.

Field-based ethnography treats access to informants embedded in communities as its most precious resource [10]. As the discipline expanded in topical scope and methodology, it retained its focus on extracting meaningful information by seeking analytical depth through engagements with relatively small numbers of informants [9]. This depth is typically achieved in part with long, repeated interviews with informants. Researchers then transcribe the text and associate transcripts to keywords, called *codes*, which form an ontology of concepts relevant to describe the problem at hand. These codes emerge from the ethnographer's embeddedness in the community she studies, gleaned through extended participant-observation which contextualises interview data in informants' larger environment [8,11].

Confronted with the rise of the global Internet, qualitative social scientists followed two main paths. One consists in mining social networking websites and applying quantitative analysis techniques to the retrieved material [12,15]. We do not discuss it here.

The other consists in convening online conversations specifically to debate the issue at hand, and treating those conversations as ethnographic data. In this approach, informants co-construct and sustain visible themes of conversation through interaction with the researcher. Further, when an ethnographer is synchronically doing research with informants, she can contextualise the temporal unfolding of information rather than getting lost in noise as in other methods that analyse aggregated digital data after the fact [5]. This approach relates to works such as participant-observation with UNIX user-groups [13], online research with Anonymous hackers [6], and fieldwork in virtual worlds like World of Warcraft [16] and Second Life [3]. In these studies, anthropologists conducted long-term ethnography, interacting with informants in-setting, asking questions, and generating context-specific data that evolved through interactions with informants over time. Some projects included offline components [13], while others were completely undertaken online [3], but all pay close attention to the ways informants make sense of their own worlds and define their terminology.

To process ethnographic evidence at scale, we recast it as *data*. Data are characterised by a structure, common to all datapoints in a given dataset, that makes it amenable to being processed by machines. Machine processing, in turn, paves the way to research at scale. The specific challenge for ethnographers is to fit their evidence into a data structure without compromising its rich, contextual character. In this section, we describe the data structure we implemented in the course of a project called OpenCare.[1] It explores how communities of any kind provide health and social care, when neither states nor business can or will serve them. It consists mostly of an online interaction environment, where individuals share their experiences of care with others, discuss, and compare notes.

2.1 Primary Data: Contributions

SSN-based ethnographies start with the posts/comments on the social networking platform or online forums that hosts the conversation. We call *contribution* a testimony in written form (interview transcript, post on an online forum, etc.). A contribution is a datapoint of the study's primary dataset, the one generated by the informants themselves. A minimum viable structure for encoding a contribution as primary data includes:

Contribution ID The contribution's unique identifier.
Text The contribution's complete text.
Author ID A unique identifier for the informant that contributed the text.
Target ID A unique identifier for the informant that the text is addressed to.
Date and time

[1] http://opencare.cc.

2.2 Secondary Data: Annotations

As we noted in Sect. 1, ethnographers work by associating snippets of texts in the primary data to codes (keywords). This generates an ontology representative of the corpus of evidence at hand. In doing so, researchers produce secondary data. We call *annotation* the atomic result of this activity. A minimum viable structure for annotations as secondary data includes:

Annotation ID The annotation's unique identifier.
Contribution ID The unique identifier of the post or comment that this annotation refers to.
Snippet The part of the text in the contribution that the researcher wishes to associate with the code.
Code The ethnographic code associated to the snippet.
Author ID Unique identifier for the researcher that produced the annotation. It is useful in the case of multi-author studies.
Date and time.

This representation is sufficient to induce a network where the nodes are informants, and edges represent interactions. Codes - associated to the interaction via annotations - encode the semantics of that interaction. This is what we call a SSN (Fig. 1). It proved to be easy to implement with most forum or blogging software applications; simple to process in meaningful ways (see below); and scalable. We propose it is general enough to fit the evidence from most ethnographies, while still rigid enough to encode it into well-formed datasets.

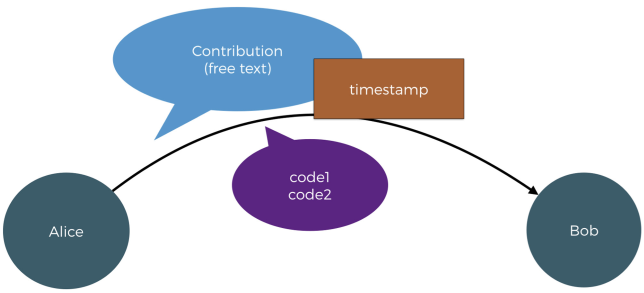

Fig. 1. An interaction between two informants that carries one or more ethnographic codes: the atom of SSNs.

3 An Application: The OpenCare Data

The OpenCare project explores how communities of any kind provide health and social care, when neither states nor businesses can or will serve them. Data are gathered from an online forum where individuals share and discuss their experiences of community-provided care. At the time of writing, the forum consists of

439 long-form posts with 2,082 comments, authored by 254 unique informants. These were uploaded onto the online forum in the period between January 2016 and May 2017. This corpus was enriched with 4,555 annotations, employing 1,035 unique codes.

3.1 The OpenCare Social Network

Online conversation in OpenCare induces a social network where nodes are community members, and edges encode interaction. For two users A and B, we produce a connection $A \rightarrow B$ if A has commented B's content at least once. This network is directed ($A \rightarrow B \neq B \rightarrow A$) and weighted (the edge $A \rightarrow B$ has a weight of k if A has commented B's content k times). It has 249 nodes and 1,007 edges The main feature of this network is that there are no signs of polarisation, nor of balkanisation. Almost all participants are connected to the giant component, so that information is allowed to flow freely across the network. The giant component itself is not obviously resolved into distinct sub-communities (the Louvain modularity value is 0.33) [2]. These structural features allow us to tentatively accept that most opinions expressed in the forum have been run past someone (other than the proponent) in the conversation.

3.2 The OpenCare Semantic Network

The network representation that proved the most useful to ethnographic research is what we call the *co-occurrence network*. Its nodes are codes. Its edges are induced between two codes that occur in annotations that refer to the same post or content. This network is undirected ($a \rightarrow b \equiv b \rightarrow a$) and weighted (the edge has a weight of k if a co-occurs with b on k different posts or comments).

We can think of the co-occurrence network as an association map between the concepts expressed by the codes. A higher edge weight k indicates a stronger group-level association between the two codes connected by the edge.

The annotations on the OpenCare corpus induce a co-occurrence network with 1,035 nodes and 12,785 edges. The main component is formed of 990 nodes and 12,777 edges, and is a perfect example of small-world network as defined in [20], with a high average clustering coefficient $\bar{C} = 0.711$.

4 Results and Discussion

4.1 Filtering the Co-occurrence Network for a High-Level View

We can think of the codes co-occurrence network as a concept association map. Rather than representing the point of view of a single individual, it encodes contributions from all informants, since informants are known to be in conversation with each other about the topic at hand. The resulting concept map, therefore, does not simply *aggregate* the association patterns of individuals, like a survey; it is the product of the *interaction* across participants. Edge weight k, then, represents the strength with which the conversation as a whole (as opposed to its participants) associates the codes connected by that edge.

Filtering the graph by higher value of k allows the researcher to single out the strongest associations between codes made by the informants as a group. Figure 2 shows the OpenCare code co-occurrence network for $k \geq 5$ (69 codes, 96 edges). The color coding was generated by applying to it the Louvain community detection algorithm [2]. It is a highly modular network (modularity = 0.59), with clearly distinguishable groups of codes. This suggests that the debate self-organised into sub-issues.

Even with no access to the primary data, one can glimpse the structure of the debate among informants. Consider for example the cluster around `legality` (in blue): we find `existing system failure` and `regulation`, reflecting the preoccupation of some informants that community health care initiatives (much needed when systems fail), turn out to be illegal. We also find `safety`, reflecting the acknowledgement that regulation is there for a reason.

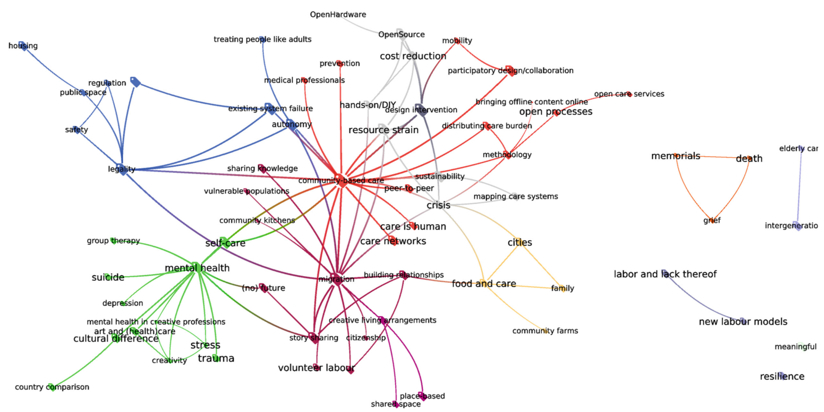

Fig. 2. The OpenCare code co-occurrence network (filtered for $k \geq 5$) (Color figure online).

4.2 Exploring the Co-occurrence Network for Novelty

Connections that spark interest can be explored at a more granular level, either by focusing on lower-k co-occurrences between a chosen code and others (not shown), or by reading the original stories and posts in which the co-occurrences took place. Around `mental health`, `stress` and `trauma`, one finds the specific (often low-cost and social) tools that people use to manage stress, for example, learning about the usefulness of `group therapy` and online support groups, alternative therapies like `holistic healthcare`, or simply `getting outdoors`. The method allows for rich analysis on multiple levels, and never loses the granularity that makes ethnographic research so powerful.

Even at high levels, meaning comes through clearly and, crucially, high level connections illuminate connections that would be invisible at a smaller scale of analysis (for example the high-k connections between `mental health` and `trauma`, and `creativity`, `art and health care` and `story sharing`). In OpenCare, for example, these high-level connections made visible by the co-occurrence graph have enabled us to theorise that people, faced with health care issues, turn to technology to connect them to each other, and we have been able to detail and verify that theory through engagement with the granular details present in the stories. Without the co-occurrence network, vital interconnections would have been missed; without the detailed ethnographic data, the meaning behind those connections would be lost.

5 Extensions and Future Research

5.1 Open Data and Large-Scale Collaboration in Ethnography

Social Semantic Networks hold the potential to make ethnography a large-scale collaborative discipline. For this to happen, we propose ethnographers embrace the practice of using and publishing open data. Open data are data that are (a) machine-readable, (b) published under licenses that allow their re-use, and (c) documented with appropriate metadata[2]. This paves the way for:

1. *Replication.* An ethnographer could pull in a colleague's primary and secondary data and check that the latter's process is clear. This increases the accountability of the research process.
2. *Large scale studies.* Accurate documentation of the code ontology allows ethnographers to work consistently on projects that would be too large for a single ethnographer to tackle. This would allow ethnographic studies at the scale of the thousands of informants.
3. *Multilingual studies.* The code ontology can be structured as a hierarchy, so that codes with the same meaning in different languages are entered in the secondary data as children of the same parent code. For example, `labour` could have `travail` and `arbeit` as children. The code co-occurrence network would be drawn between parent nodes, thus allowing both an all-languages view on data and across-languages comparisons.
4. *Reuse and extension.* An ethnographer could pull in a colleague's primary and secondary data, add her own coded corpus and use the combination of annotated corpora to produce a completely new study - for example on responding to the refugee crisis, one of the care issues taken up by the OpenCare community.
5. *"Longitudinal ethnography"*. An online conversation could be revamped over time (for example every year) to keep track of how its collective point of view evolves.

[2] We have done this with OpenCare: https://doi.org/10.5281/zenodo.164970; https://github.com/opencarecc/opencare-data-documentation.

The combination of these elements requires a cultural shift from practitioners. Ethnographers tend to work alone, and there is as yet no culture of open data in the discipline, as access to coded interviews and fieldnotes belongs to the ethnographer alone. To the best of our knowledge, the OpenCare dataset is the first-ever open dataset including primary and secondary data.

Yet, the payoff of such a shift is substantial. Ethnography could bring its unique methodological advantages to new problems, that demand a scale and consistency it now cannot supply. For example, we could imagine a version of Eurobarometer based on an open online conversation. Instead of answering multiple choice questions (vulnerable to framing biases [19]), informants would discuss their perception of Europe, allowing researchers to discover novel patterns of association and detect the fading of old ones.

5.2 Other Methodological Improvements

In the future, we plan to test with at least three improvements:

1. Weighing contributions (and consequently annotations) by a "reliability score" derived by applying social theory on the social network topology [4].
2. Applying alternative ways to measure edge (association) strength k in the co-occurrence network; for example, $k(a \rightarrow b)$ could encode the number of informants that have authored contributions coded with both codes a and b, or the number of separate threads which contain at least one contribution with it, and so on. Different measures of edge strength have different interpretations, so they would allow different perspectives on the data corpus.
3. Observing and modelling the online conversation as a dynamic system. Stochastic Actor-Oriented Models might be a good place to start, despite known limitations [18].

6 Conclusions

Semantic social networks show some promise as a method for social research aimed at capturing collective intelligence [14], with some of the advantages of both purely qualitative traditional ethnography and quantitative surveys. Like traditional ethnography (but unlike surveys), they deal well with open questions and novelty. Like surveys (but unlike traditional ethnography), they can handle hundreds of informants. When combined with open standards and open data, they could perhaps attempt to handle thousands of informants. We look forward to exploring this potential further.

Acknowledgment. The OpenCare project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement no 688670).

References

1. Barben, D., Fisher, E., Selin, C., Guston, D.H.: 38 anticipatory governance of nanotechnology: foresight, engagement, and integration. In: The Handbook of Science and Technology Studies, p. 979 (2008)
2. Blondel, V.D., Guillaume, J.-L., Lambiotte, R., Lefebvre, E.: Fast unfolding of communities in large networks. J. Stat. Mech. Theor. Exp. **2008**(10), P10008 (2008)
3. Boellstorff, T.: Coming of Age in Second Life: An Anthropologist Explores the Virtually Human. Princeton University Press, Princeton (2008)
4. Burt, R.S.: Brokerage and Closure: An Introduction to Social Capital. Clarendon Lectures in Management Studies. OUP, Oxford (2005)
5. Coleman, G.: Ethnographic approaches to digital media. Annu. Rev. Anthropol. **39**(1), 487–505 (2010)
6. Coleman, G.: Hacker, Hoaxer, Whistleblower, Spy: The Many Faces of Anonymous. Verso, Brooklyn (2015)
7. Conley, J.M., Williams, C.A.: Engage, embed, and embellish: theory versus practice in the corporate social responsibility movement. J. Corp. Law **31**, 1 (2005)
8. Emerson, R., Fretz, R., Shaw, L.: Writing Ethnographic Fieldnotes. University of Chicago Press, Chicago (2011)
9. Fox, R.G. (ed.): Recapturing Anthropology: Working in the Present. School of American Research Press, Santa Fe (1991)
10. Geertz, C.: The Interpretation of Cultures. Basic Books, New York City (1973)
11. Goffman, E.: On fieldwork. J. Contemp. Ethnogr. **123**(18), 123–132 (1989)
12. Hollstein, B.: Qualitative Approaches, pp. 404–416. SAGE Publications Ltd., Thousand Oaks (2011)
13. Kelty, C.: Two Bits The Cultural Significance of Free Software. Duke University Press, Durham (2008)
14. Lévy, P.: Collective Intelligence: Mankind's Emerging World in Cyberspace. Plenum Trade, New York (1997)
15. Munk, A.K., Abildgaard, M.S., Birkbak, A., Petersen, M.K.: (Re-)appropriating instagram for social research: three methods for studying obesogenic environments. In: Proceedings of the 7th 2016 International Conference on Social Media & Society, SMSociety 2016, pp. 19:1–19:10. ACM (2016)
16. Nardi, B.: My Life as a Night Elf Priest: An Anthropological Account of World of Warcraft. University of Michigan Press, Ann Arbor (2010)
17. Rheingold, H.: The Virtual Community: Homesteading on the Electronic Frontier. MIT Press, Cambridge (2000)
18. Snijders, T.A.B.: Stochastic actor-oriented models for network change. J. Math. Sociol. **21**(1–2), 149–172 (1996)
19. Tversky, A., Kahneman, D.: The framing of decisions and the psychology of choice. In: Covello, V.T., Mumpower, J.L., Stallen, P.J.M., Uppuluri, V.R.R. (eds.) Environmental Impact assessment, Technology Assessment, and Risk Analysis. NATO ASI Series, vol. 4, pp. 107–129. Springer, Heidelberg (1985). doi:10.1007/978-3-642-70634-9_6
20. Watts, D.J., Strogatz, S.H.: Collective dynamics of 'small-world' networks. Nature **393**, 440–442 (1998)
21. Wolford, W.: From confusion to common sense: using political ethnography to understand social mobilization in the Brazilian Northeast. In: Joseph, L., Mahler, M., Auyero, J. (eds.) New Perspectives in Political Ethnography, pp. 14–36. Springer, Heidelberg (2007). doi:10.1007/978-0-387-72594-9_1

A Topic Detection and Visualisation System on Social Media Posts

Stelios Andreadis(✉), Ilias Gialampoukidis, Stefanos Vrochidis, and Ioannis Kompatsiaris

Information Technologies Institute, Centre for Research and Technology Hellas, Thessaloniki, Greece
{andreasisst,heliasgj,stefanos,ikom}@iti.gr
http://mklab.iti.gr/

Abstract. Large amounts of social media posts are produced on a daily basis and monitoring all of them is a challenging task. In this direction we demonstrate a topic detection and visualisation tool in Twitter data, which filters Twitter posts by topic or keyword, in two different languages; German and Turkish. The system is based on state-of-the-art news clustering methods and the tool has been created to handle streams of recent news information in a fast and user-friendly way. The user interface and user-system interaction examples are presented in detail.

Keywords: Topic detection and visualisation · Twitter posts · Keyword-based search · Topic-based filtering

1 Introduction

The increasing number of social media users ranges from young people to the elderly people, having the need to identify and detect interesting topics and events. This is a challenging task due to the large amount of posts that are published on a daily basis. Topic detection aims at grouping together social media posts or text documents in general that discuss about the same topic-event. Topics are not considered general categories, such as politics, sports or lifestyle, but particular thematic areas and trending topics which are updated. Topic labels also need regular updates, according to the most recent documents that are collected. Fields of application for topic detection vary from media monitoring and news recommendation to content linking in the security domain.

In the social media domain, several monitoring tools have appeared, offering search by keyword, statistics and management of multiple accounts, such as Hootsuite[1] and TweetReach[2]. In addition, Social Mention[3] displays top keywords, hashtags, and sites and Twazzup[4] shows the top influencers (based on

[1] https://hootsuite.com/.
[2] https://tweetreach.com/.
[3] http://socialmention.com/.
[4] http://new.twazzup.com/.

I. Kompatsiaris et al. (Eds.): INSCI 2017, LNCS 10673, pp. 421–427, 2017.
https://doi.org/10.1007/978-3-319-70284-1_33

followers), the most retweeted photos, and the top 10 keywords in response to a search by query. Other tools that allow the social media user to handle efficiently large streams of posts include IceRocket, TweetDeck, Twitonomy, Followerwonk and SumAll. Contrary to these approaches, our tool is able to cluster a selection of recent social media posts by topic, as identified from a set of concepts that are assigned to each post. Each detected topic is visualised using a cloud of concepts and named entities, as identified in the collection. Moreover, the tool offers a language option and is able to filter results by a list of suggested keywords. The tool has been developed for the purposes of the H2020-KRISTINA[5] project, building on top of other relevant topic detection services (e.g. FP7-MULTISENSOR[6]), hence the supported languages are German and Turkish.

Topic detection assumes a vector representation of a text document and is usually considered as a clustering problem [1], in absence of training sets. One of the most popular topic modeling methods is Latent Dirichlet Allocation (LDA), which requires as input the number of topics. On the other hand, density based approaches [2–4] do not require a priori knowledge of the number of clusters, but they are less effective than LDA in text clustering. Moreover, LDA has been generalised to nonparametric Bayesian approaches, such as the hierarchical Dirichlet process [5] and DP-means [6]. The estimation of the correct number of topics has been examined in [7], where the DBSCAN-Martingale has shown better performance than other state-of-the-art methods [8]. Other topic detection approaches in the social media domain involve graph-based approaches [9], where a graph clustering algorithm is applied on the graph of text documents and the decision, whether to link two social media posts or not, is based on the output of a classifier, which assigns or not the candidate items in the same cluster. Contrary to this graph-based approach, we cluster social media posts in an unsupervised way, using the DBSCAN-Martingale to estimate the number of clusters and then LDA to assign posts to topics.

Our paper is structured as follows. In Sect. 2 we describe the overall system architecture and we refer to the methods which are involved in the data analysis. In Sect. 3 we present the user interface of the topic detection tool, as well as user interaction modes with concrete usage scenarios.

2 Topic Detection System

The overall system architecture of our proposed tool is presented in Fig. 1. The front page offers search by keyword as a database query. Our MongoDB database currently contains 48,589 tweets of 8 well-known journalism organizations in Germany and Turkey, crawled from May 2, 2017 up to now by exploiting the Twitter API[7]. Topic detection is a service, developed in Java, that builds on top of open source code[8], which is based on the combination of density-based

[5] http://kristina-project.eu/en/.
[6] https://www.multisensorproject.eu/.
[7] https://dev.twitter.com/.
[8] https://github.com/MKLab-ITI/topic-detection.

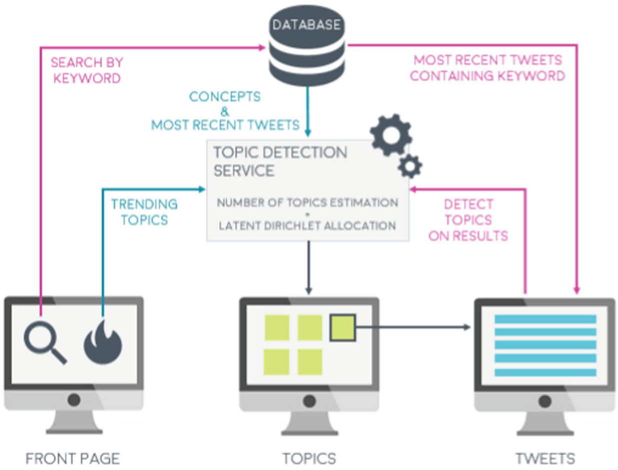

Fig. 1. System architecture

clustering with LDA, as proposed in [7]. The module estimates the number of clusters (topics) and the estimation is followed by Latent Dirichlet Allocation [10], so as to assign social media posts to topics. Topic detection takes place either on the most recent (using a timeframe option) posts or on the ones that are provided by a search by keyword. The second approach has also been presented in [11], as query-based topic detection. Contrary to these approaches, the method is more efficient, since it has been recently observed that a faster version of the so called "DBSCAN-Martingale" [7] appears when sampling the density levels of DBSCAN from a skewed distribution rather than the uniform distribution [12].

3 Topic Detection Interface and Interaction Modes

Our Topic Detection Interface[9] has been designed to offer the end user an efficient and intuitive way to search for recent tweets of interest and to automatically identify topics in the resulted content. The application was implemented using common Web technologies, i.e. HTML5, CSS, JavaScript, jQuery, and PHP to communicate with the database, as well as the Kendo UI Core framework[10].

3.1 Topic Detection Interface

The front page, as depicted in Fig. 2, includes a one-line text input field, common to the user's experience of popular search engines. The user is able to type a set of keywords and then initiate a search for tweets that contain one or more of these keywords. Pressing Enter on keyboard or clicking the magnifier icon

[9] http://mklab-services.iti.gr/KRISTINA_topic_detection/.
[10] https://github.com/telerik/kendo-ui-core.

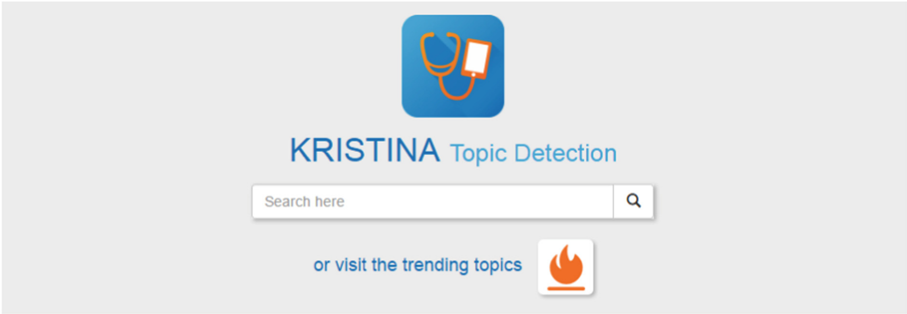

Fig. 2. Front page of the user interface

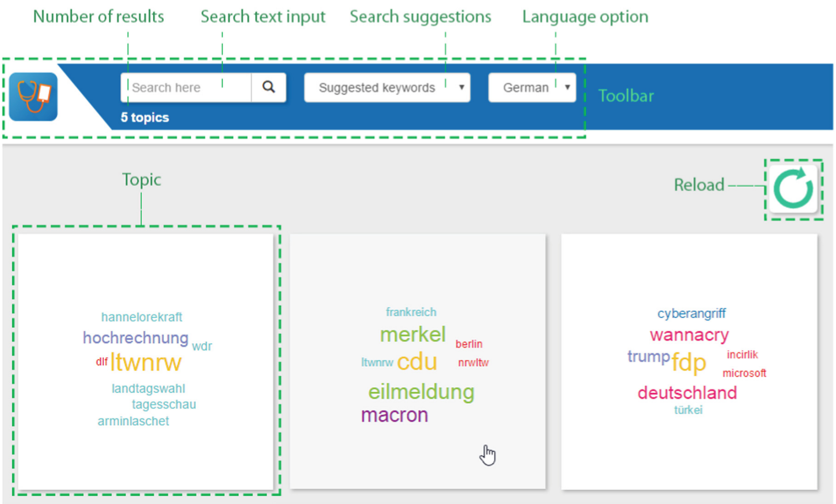

Fig. 3. Screenshot of detected topics in German; from left to right they refer to the elections in North Rhine-Westphalia (NRW), the meeting of Merkel and Macron, and the WannaCry ransomware attack.

will navigate to the results page. Alternatively, by clicking the flame button under the search input field, the user can visit the trending topics, which will be described later in this section.

The results page, shown in Figs. 3 and 4, consists of two basic components; a toolbar on top of the user interface that can be seen in detail in Fig. 3 and the main panel that covers most of the page. Describing the toolbar from left to right, there is a logo that returns to the front page, a text input field to facilitate more searches and the number of results, either tweets or topics. Moreover, the toolbar includes a dropdown selection of keywords that will be explained along with the topic detection and a dropdown module to select language. The default language is German, but the user also has the option to perform search or topic detection in Turkish.

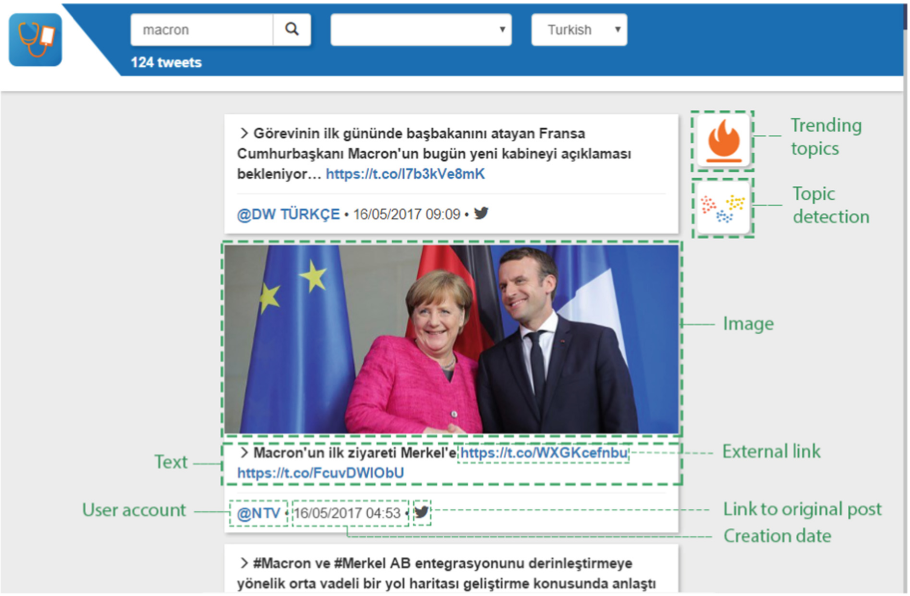

Fig. 4. Screenshot of resulted tweets in Turkish about Macron's visit to Germany

Regarding the main panel, the outcome of the search is shown as a list of white boxes, where every box represents a tweet (Fig. 4). Please note here that the user interface connects directly to the database and finds the five hundred most recent posts that include any of the keywords defined by the user. For each tweet, a variety of information is presented. In detail, the main text of the post, together with images or active links if existing, the username of the author, which links to the user account's Twitter page, the date when the tweet was published, and a link to the original post in Twitter (by clicking on the grey Twitter icon). All results are sorted by their creation date, from most recent tweet to oldest.

After search results are displayed, the user is provided with two different options, in the form of two buttons on the right. The first button with the flame icon navigates to the trending topics, while the second button performs topic detection on the resulted tweets. When the latter is clicked, the topic detection service is called, receiving the ids of the tweets as input and returns the topics detected. Then, tweets are removed and the main panel is filled with topics in a grid-like view (Fig. 3). Each topic is represented as a square, white box with its most frequent concepts or named entities, as extracted in the context of the H2020-KRISTINA project, in a colorful word cloud, where the more recurring terms appear in larger font. These terms are also added to the aforementioned dropdown selection in the toolbar, serving as suggested keywords to start a new search. The reload button on the top right corner can be used to call again the service and get new topics, since detection is performed dynamically. Clicking on a topic box reveals the tweets it consists of, in a manner similar to search results.

Finally, as it has already been mentioned above, another module of the user interface is the trending topics. While currently available Twitter trends refer to the most popular topics/hashtags at this given moment without further exploration criteria, the proposed system provides much more room for investigating trends by offering the option to select a target period of time, e.g. the last two days, source streams, such as interesting media and user accounts, and language, useful for trend localization. To perform this topic detection method, a variation of the previous service is invoked, passing no arguments as input (formerly set of tweet ids), but rather pre-configuring filtering options, i.e. language, user accounts and time period, in the system. The trending topics are shown exactly like the topics in Fig. 3 and clicking on them navigates to the list of tweets they are composed of, displayed in the same way as previously described, e.g. in Fig. 4.

3.2 Interaction Modes

To illustrate the functionality of the proposed interface, we provide a simple usage scenario. Supposing that the user is not interested in a particular topic and desires to be informed about recent news in general, he/she can begin by clicking the trending topics button in the front page of Fig. 2. A set of most talked-about topics on Twitter appears as in Fig. 3 and the user finds interesting the visit of French president Macron to Germany[11]. After clicking on this topic, he/she is able to read all the relative tweets and follow external links for more information. In order to find out more about Macron, the user selects the term "Macron" from the dropdown list of suggested keywords and initiates a search. All tweets in database containing the name of the French president are returned and can be read. The user is curious about the popularity of Macron in Turkey, so he/she changes the language by the available option in the toolbar and clicks the magnifier icon to rerun the search. When new tweets are shown in Turkish (Fig. 4), the user prefers to separate them in topics before reading all of them, thus he/she clicks the topic detection button and is presented with topics based on the previous results. A click on the topic with the word "almanya" ("Germany" in Turkish) provides the user with tweets about Macron's visit to Berlin, as posted by Turkish accounts.

4 Conclusion

We have demonstrated a topic detection and visualisation tool, that filters Twitter posts by topic or keyword, in two different languages; German and Turkish. The system offers to the user multiple navigation options to handle large streams of recent news information. The suggested keywords also help the user in his/her attempt to search for content, according to his/her interests or needs. In the future, we plan to advance this system using additional languages and other functionalities for a more user-friendly experience. Furthermore, exploiting user feedback to automatically create and update a training corpus could lead to more personalised information.

[11] http://edition.cnn.com/2017/05/15/world/macron-merkel-meeting/.

Acknowledgements. This work was supported by the EC-funded projects H2020-645012 (KRISTINA) and H2020-700475 (beAWARE).

References

1. Aggarwal, C.C., Zhai, C.: A survey of text clustering algorithms. In: Aggarwal, C., Zhai, C. (eds.) Mining Text Data, pp. 77–128. Springer, Heidelberg (2012). doi:10.1007/978-1-4614-3223-4_4
2. Ester, M., Kriegel, H.P., Sander, J., Xu, X.: A density-based algorithm for discovering clusters in large spatial databases with noise. In: Proceedings of the 2nd International Conference on Knowledge Discovery and Data Mining (KDD 1996), pp. 226–231. AAAI Press (1996)
3. Ankerst, M., Breunig, M.M., Kriegel, H., Sander, J.: OPTICS: ordering points to identify the clustering structure. In: Proceedings of the ACM SIGMOD International Conference on Management of Data, pp. 49–60. ACM (1999)
4. Campello, R.J.G.B., Moulavi, D., Sander, J.: Density-based clustering based on hierarchical density estimates. In: Pei, J., Tseng, V.S., Cao, L., Motoda, H., Xu, G. (eds.) PAKDD 2013. LNCS, vol. 7819, pp. 160–172. Springer, Heidelberg (2013). doi:10.1007/978-3-642-37456-2_14
5. Teh, Y.W., Jordan, M.I., Beal, M.J., Blei, D.M.: Sharing clusters among related groups: Hierarchical Dirichlet processes. In: NIPS, pp. 1385–1392 (2004)
6. Kulis, B., Jordan, M.I.: Revisiting k-means: new algorithms via Bayesian nonparametrics. arXiv preprint arXiv:1111.0352 (2011)
7. Gialampoukidis, I., Vrochidis, S., Kompatsiaris, I.: A hybrid framework for news clustering based on the DBSCAN-martingale and LDA. Machine Learning and Data Mining in Pattern Recognition. LNCS, vol. 9729, pp. 170–184. Springer, Cham (2016). doi:10.1007/978-3-319-41920-6_13
8. Charrad, M., Ghazzali, N., Boiteau, V., Niknafs, A., Charrad, M.M.: Package 'NbClust'. J. Stat. Softw. **61**, 1–36 (2014)
9. Petkos, G., Schinas, M., Papadopoulos, S., Kompatsiaris, Y.: Graph-based multimodal clustering for social multimedia. Multimedia Tools Appl. **76**, 7897–7919 (2017). doi:10.1007/s11042-016-3378-2
10. Blei, D.M., Ng, A.Y., Jordan, M.I.: Latent dirichlet allocation. J. Mach. Learn. Res. **3**(Jan), 993–1022 (2003)
11. Gialampoukidis, I., Liparas, D., Vrochidis, S., Kompatsiaris, I.: Query-based topic detection using concepts and named entities. In: Proceedings of the 1st International Workshop on Multimodal Media Data Analytics (MMDA 2016), pp. 9–13 (2016)
12. Gialampoukidis, I., Vrochidis, S., Kompatsiaris, I., Antoniou, I.: Topic detection using the DBSCAN-Martingale and the time operator. In: 17th Conference of the Applied Stochastic Models and Data Analysis (ASMDA2017) (2017)

Towards Suicide Prevention: Early Detection of Depression on Social Media

Victor Leiva and Ana Freire(✉)

Department of Communications and Information Technologies,
Universitat Pompeu Fabra, Carrer Tanger, 122-140, 08018 Barcelona, Spain
victor.leiva01@estudiant.upf.edu, ana.freire@upf.edu

Abstract. The statistics presented by the World Health Organization inform that 90% of the suicides can be attributed to mental illnesses in high-income countries. Besides, previous studies concluded that people with mental illnesses tend to reveal their mental condition on social media, as a way of relief. Thus, the main objective of this work is the analysis of the messages that a user posts online, sequentially through a time period, and detect as soon as possible if this user is at risk of depression. This paper is a preliminary attempt to minimize measures that penalize the delay in detecting positive cases. Our experiments underline the importance of an exhaustive sentiment analysis and a combination of learning algorithms to detect early symptoms of depression.

Keywords: Early detection · Depression · Social media · Machine learning

1 Introduction

Suicide has become an enormous public health problem worldwide. According to the World Health Organization (WHO), approximately 804,000 deaths occurred due to suicide in 2012, becoming the first cause of death between the ages of 15 and 44. It is worth highlighting that in high-income countries the amount of young adults (i.e., between 15 and 29 years old) who commit suicide accounts for 17.6% of the total number of deaths [13]. The statistics presented by the WHO inform that 90% of the suicides can be attributed to mental illnesses in high-income countries [12]. Besides, Park et al. [9] concluded that people with mental illnesses tend to reveal their mental condition on the social media, as a way of relief. Among all users of these social media platforms, adolescents have been pinpointed as the most frequent ones [3].

Hence, these previous studies drive us through the detection of depression on social media as a first step against suicidal behaviour.

Using a labeled dataset [7] with messages posted in *Reddit*[1], we experiment different techniques for improving previous results for early detection of depression considering time-aware classification measures.

[1] Reddit: https://www.reddit.com.

I. Kompatsiaris et al. (Eds.): INSCI 2017, LNCS 10673, pp. 428–436, 2017.
https://doi.org/10.1007/978-3-319-70284-1_34

2 Background

In the last few years, more research about detecting mental illnesses on social media has become available [1,2,4,7,8,11,14]. Twitter has been widely investigated [2,4,8], although other social networks have also been discussed in the literature, such as Sina [11] or Myspace [1].

De Choudhury et al. [4] proposed a probabilistic model to determine which tweets could indicate depression. They extract both post-based (i.e. positive or negative effect, linguistic style) and user-based features (i.e. number of postings, number of followers/followees). After applying dimensionality reduction (Principal Component Analysis - PCA) and using Support Vector Machine (SVM) they got 74% accuracy.

A more recent work from 2016 [8] also examined Twitter with the same purpose. In this case, they use a bag of words (counts the word occurrence frequencies to quantify the content of a tweet) to vectorize the tweets. Then, they applied different machine learning algorithms (Logistic Regression, Naïve Bayes and SVM) and obtained around 80% accuracy in detecting depressive disorder. Logistic Regression has also been applied with similar objective in [1], in this case to classify Myspace post-mortem comments exhibiting (or not) emotional distress. Authors made use of Linguistic Inquire and Word Count library (LIWC)[2] to perform text analysis and get its polarity.

In the same context, sentiment analysis using LIWC was also applied to create a depression detection model for Chinese micro-blogs [11]. This model gathered 10 features from the text (again, the polarity of the text but also other innovative features such as the use of emoticons). The initial 10 features were later reduced to 5, with a corresponding degradation in the results but a high improvement regarding computational time and amount of data needed.

Lately, *Twitter* has become the most popular network among researchers. Nevertheless, a recent work by Losada and Crestani et al. [7] draw our attention to *Reddit*. Authors argued that the limitation in the number of characters in a tweet reduces the context we can get about the writer. Reddit has no limitation regarding the number of characters and we can get unlimited access to previous submissions from the same user, as opposite to Twitter, that only allows to download up to 3200 tweets per user. Authors created and distributed a dataset including *Reddit* messages from different users, some of them diagnosed with depression. Such dataset was labeled indicating if the users were at risk of depression. Losada et al. noted that the most widely used classification measures (precision, recall, F-measure) are time-unaware, so they do not reward early alerts. Therefore, they proposed a new metric for early detection that penalizes the delay in detecting positive cases. They also reported the results of some preliminary baselines to detect early symptoms of depression. These baselines are quite simple and the features extracted from the text are just based on the tf-idf vectorization.

[2] LIWC: http://liwc.wpengine.com/.

3 Proposal

We aim to improve the state of art of time-aware early detection of depression on social media by means of:

- Considering innovative measures that are specific for early detection.
- Extracting text features in a more accurate way than the current approaches (just based on tf-idf vectorization), by including sentiment analysis.
- Exploring the behaviour of classic and modern machine learning techniques to better predict positive cases.
- Studying if genetic algorithms can improve the previous results.

We give more details about this proposal in the next subsections.

3.1 Feature Extraction

We run our experiments using the dataset built by Losada and Crestani et al. [7]. It consists of a list of users catalogued as depressive (135 of them) and a control group (752 users), each of them with a corresponding random number of messages (there are users with 10 messages up to 2000, with an average of 607 messages per user). These messages cover different topics, as the redditors are often active in different subreddits. The text collection contains messages from a wide period of time so we can observe the language evolution used by a depressive subject and a non-depressive one.

In addition to tf-idf vectorization (12,968 features), we have introduced three more features that are related to the sentiment polarization of the text. The three features correspond to the positive, neutral and negative percentages associated to each message/user. To implement this sentiment analysis we employed VADER Sentiment Analysis (Valence Aware Dictionary and sEntiment Reasoner) [5]. It consists of a lexicon and a rule-based sentiment analysis tool suitable for sentiments expressed in social media.

3.2 Learning Algorithms

We have used four prediction methods. These methods are briefly described below.

- *Logistic Regression*: calculates the probability of a categorical variable (depressed/non-depressed) from a number of predicting variables.
- *Support Vector Machine*: finds a hyperplane that divides two categorical data by a clear gap that is as wide as possible.
- *K-Nearest Neighbor*: is a non-parametric method. An object is classified by a majority vote of its k neighbors. The value of each vote is weighted by the distance of the neighbors to the object.
- *Random Forest*: A random forest is a meta estimator that fits a number of decision tree classifiers on various sub-samples of the dataset and use averaging to improve the predictive accuracy and control over-fitting.

Additionally, we have built an ensemble model that uses all of the previous algorithms to determine the class of the subject. We have used a technique called *soft-voting* where the model determines the probability of one class with the mean of all the probabilities given by the algorithms.

3.3 Parameters Optimization

We have followed two strategies to determine which parameters are the optimal for each algorithm. Fist, we have applied grid search to optimize the algorithms independently. Then, by means of genetic algorithms, we have optimized their contribution in the ensemble method. For the second step we have chosen genetic algorithms, as they have demonstrated to be more stable and faster than grid search for the optimization of high number of parameters [6].

3.4 Evaluation

For results evaluation, we have used the measures proposed by Losada and Crestani et al. [7]: Early Risk Detection Score *ERDS*, defined as follows:

$$ERDS_o(d,k) = \begin{cases} c_{fp} & \text{if d = false positive} \\ c_{fn} & \text{if d = false negative} \\ lc_o(k) \cdot c_{tp} & \text{if d = true positive} \\ 0 & \text{if d = true negative} \end{cases} \tag{1}$$

where $c_{fp} = 0.21$, $c_{fn} = 1$, $c_{tp} = 1$ and the function $lc_o(k)$ has the next structure:

$$lc_o(k) = 1 - \frac{1}{1 + e^{k-o}} \tag{2}$$

All the scores are normalized by the number of testing users.

We also report the traditional score functions like the precision, recall and the F1 score.

4 Experimental Setup

4.1 Dataset Setup

The *Reddit* dataset is structured in XML files, one for each redditor, that contains each of her submission ordered chronologically. For each post, we have its title, its text and the publication date. We vectorize the text using tf-idf. We remove stopwords and only select the terms that appear in 20 documents or more, following Losada and Crestani et al. [7] approach.

Since this kind of vectorization leads to very sparse data we use Principal Component Analysis (PCA) in order to reduce the dimensionality of the vector to the most informative components. We apply grid search to set the dimensionality of PCA to 300 features, while initially we had 12,968 tf-idf features (see Fig. 1).

Both the implementation of the learning algorithms and the vectorization are implemented with scikit-learn library, version 0.18, for Python [10].

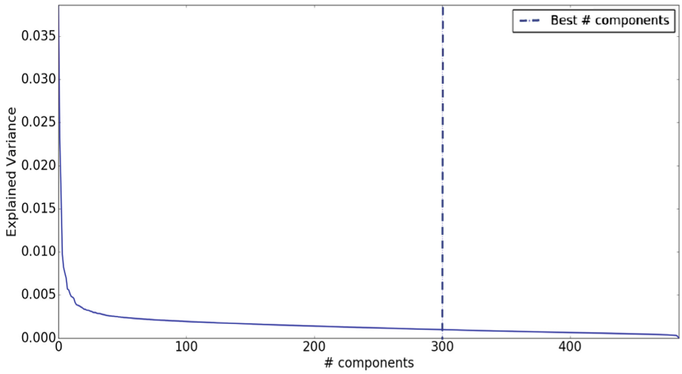

Fig. 1. Principal Component Analysis (PCA)

4.2 Baselines

We have used different baselines that obtained the best performance in [7]. These baselines are based on Linear Regression combined with the following strategies to emit a decision about the users:

- **First *n*:** This method consist in concatenating the first n messages from each user and making a prediction based on the result obtained from the learning model. If n is bigger than the total number of messages the method uses all the data available. This method has a delay of n messages (we experiment with n={10, 50, 100}).
- **Dynamic**: This method incrementally concatenates the texts of a user and only makes a decision if the depression language classifier outputs a confidence value above a given threshold. We experiment with probabilities 0.50 and 0.75, as they have shown the best performance in [7]. It is based on Logistic Regression with a particular set of parameters: *penalty*=‘l1’, *solver*=‘liblinear’, C=16 and *class_weight*={0:0.2, 1:0.8}.

4.3 Experimentation Steps

Our experiments are organized in several consecutive steps.

First, we report the results we obtained with the baselines proposed by Losada and Crestani et al. [7], as the dataset they released is slightly different from the one they used to publish the results. Note that these baselines are build upon Linear Regression with tf-idf vectorization. From this analysis we select the best baseline for the next steps. We vary the parameter n of the strategy *First n* (i.e. [10, 50, 100] instead of [100, 500]), in order to select an even more demanding baseline that emits a decision earlier and so to avoid using all the messages from some users.

Then, we compare the best baseline with the proposed algorithms (see Sect. 3.2: Support Vector Machine (SVM), K-Nearest Neighbor (KNN), Random Forest (RF) and the ensemble voting algorithm (Ensemble).

Then, we apply the Principal Component Analysis (PCA) to see the effect of dimensionality reduction and we add new features resulting from VADER Sentiment Analysis.

Finally, we report the behaviour obtained with the hybrid model that uses a genetic algorithm (GA) to optimize the ensemble method algorithm described in Sect. 3.2. In the genetic algorithm we can tune, basically, five parameters: the number of individuals, the number of generations, the tolerance, the birth/death rate and the mutation rate. In our experiments we have set the $ERDS$as our tolerance with a stopping value of 5%, the number of individuals is set to 100, the maximum number of generation is 40 and the birth/death rate corresponds to 20%. For the mutation rate we have considered three different strategies. The first one maintains a constant rate over time. Therefore, we try to avoid getting stuck in local minima of the score function. However, as a drawback, if the rate is too high, the algorithm may take infinite steps to converge to the optimal values. In our experiment with the mutation rate constant, we have tuned it to a value of 30%. For the other two strategies, we have implemented two functions that decrease over generations. We wanted to keep a high value of the mutation rate in the initial generations to diversify the individuals. At the same time, we wanted a low mutation rate in further generations in order to facilitate the convergence of the algorithm. The functions selected are an Asymptotic and a Gompertz function, with the following expressions:

$$Mutation\,rate\,functions \begin{cases} Asymptotic(x) = \frac{1}{x} \\ Gompertz(x) = 1 - 0.9e^{-6e^{-0.15x}} \end{cases} \tag{3}$$

The comparison between the two functions can be found in Fig. 2. We observe that the Asymptotic function decreases more rapidly than the Gompertz function, which stays at higher values during more generations.

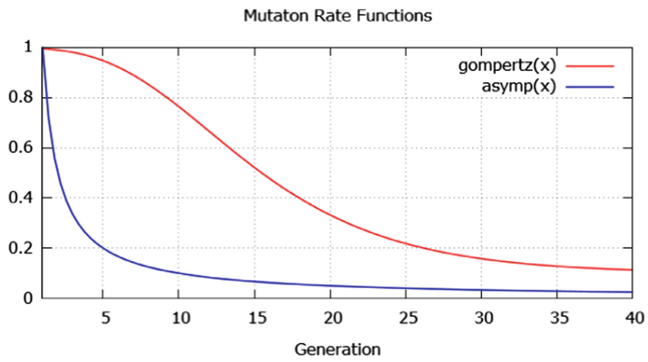

Fig. 2. Mutation rate functions.

5 Results

5.1 Baselines

The evaluation of the proposed baselines (using Linear Regression) can be found in Table 1. The results obtained from the *Dynamic* strategies obtain good results for most of the scores. The *Dynamic* strategy with probability 50% gets the best values for the early risk measures, as well as for the recall. Thus, we select linear regression with Dynamic 50% as the baseline for the next experiments.

Table 1. Baselines (linear regression)

	ERDE5	ERDE50	P	R	F1
First 10	0.10	0.09	0.63	0.29	0.39
First 50	0.08	0.08	0.57	0.46	0.51
First 100	0.07	0.07	0.55	0.56	0.55
Dynamic (prob. 50%)	**0.06**	**0.05**	0.41	**0.77**	0.54
Dynamic (prob. 75%)	0.07	0.07	0.61	0.54	**0.57**

5.2 Learning Algorithms

We report the results obtained after running the learning algorithms described in Sect. 3.2, including the dimensionality reduction (PCA) and the new features based on VADER sentiment analysis (see Table 2). LR (Linear Regression) corresponds to the baseline Dynamic (prob. 50%). We have used a grid search over a list of possible parameters combinations to select those used to train the learning algorithms.

Curiously, there are not significant differences after applying PCA and VADER, and the results highly depend on the learning algorithm. For instance, regarding LR, the combination PCA + VADER improves precision and F-measure but increases the error in case of ERDE50 and decreases recall. However, in case of KNN, after applying PCA + VADER all measures experiment some improvement, except for recall, that gets a value of 0.77 (the same as the baseline).

Although dimensionality reduction not always gets better results, it can still be applied to gain computational resources [11]. Regarding sentiment analysis, we will explore more advanced methods that also classify the text into numerous categories such as depression, happiness, sadness and so on.

Regarding the learning methods, some of them can get better results than the baseline, as we can see after applying genetic algorithms for parameter optimization with VADER (last row in the table). We run the Genetic Algorithm with the Gompertz function (highest performance) considering PCA and VADER separately. The use of the Genetic Algorithm along with the three text polarity features (VADER) performs 16,7% better than the baseline, in terms of ERDE5, while it maintains the value of ERDE50 in 0.05. Precision and F1 are also improved, and Recall is maintained at 77%.

Table 2. Proposed learning algorithms with the dynamic strategy (prob. 50%). LR is the baseline. In bold, the best result for each algorithm with and without PCA + VADER. In italic, the best result for each measure.

	ERDE5	ERDE50	P	R	F1
LR	0.06	**0.05**	0.41	**0.77**	0.54
LR + PCA + VADER	0.06	0.06	**0.49**	0.67	**0.57**
Support vector machine	0.07	0.07	*0.58*	**0.58**	**0.58**
SVM + PCA + VADER	0.07	0.07	0.58	0.56	0.57
KNN	0.07	0.07	0.23	0.88	0.37
KNN + PCA + VADER	**0.06**	**0.05**	**0.38**	0.77	**0.51**
RF	**0.08**	0.08	**0.22**	0.83	**0.35**
RF + PCA + VADER	0.09	0.08	0.19	***0.96***	0.31
Ensemble	0.06	**0.05**	0.53	**0.69**	***0.60***
Ensemble + PCA + VADER	0.06	0.06	0.53	0.67	0.59
GA (Constant) + PCA + VADER	0.08	0.07	**0.54**	0.51	0.52
GA (Asymptotic) + PCA + VADER	0.07	0.07	0.52	0.58	0.55
GA (Gompertz) + PCA + VADER	0.07	0.06	0.53	0.60	0.56
GA (Gompertz) + PCA	0.06	0.06	0.53	0.67	**0.59**
GA (Gompertz) + VADER	***0.05***	***0.05***	0.45	**0.77**	0.57

6 Conclusions

This paper has investigated how to better detect early risk of depression in social media, by optimizing time-aware classification measures: ERDE5 and ERDE50. We have applied different learning algorithms, and combinations of them. Other techniques such as dimensionality reduction and text polarity have been studied. We have provided further evidence for the benefit of applying genetic algorithms and text polarity (16,7% improvement regarding the baseline). Future work will concentrate on applying more accurate text sentiment classification to get a better representation of the input features.

Acknowledgements. This work was supported by the Spanish Ministry of Economy and Competitiveness under the Maria de Maeztu Units of Excellence Programme (MDM-2015-0502).

References

1. Brubaker, J.R., Kivran-Swaine, F., Taber, L., Hayes, G.R.: Grief-stricken in a crowd: the language of bereavement and distress in social media. In: Proceedings of ICWSM (2012)
2. Sowles, S., Connolly, S., Rosas, C., Bharadwaj, M., Bierut, L.J., Cavazos-Rehg, P.A., Krauss, M.J.: A content analysis of depression-related tweets. Comput. Hum. Behav. **1**(54), 351–357 (2016)

3. Cole, D.A., Zelkowitz, R.L., Nick, E., Martin, N.C., Roeder, K.M., Sinclair-McBride, K., Spinelli, T.: Longitudinal and incremental relation of cybervictimization to negative self-cognitions and depressive symptoms in young adolescents. J. Abnorm. Child Psychol. **44**(7), 1321–1332 (2016)
4. De Choudhury, M., Counts, S., Horvitz, E.: Social media as a measurement tool of depression in populations. In: Proceedings of the 5th Annual ACM Web Science Conference, pp. 47–56. ACM (2013)
5. Hutto, C., Gilbert, E.: Vader: a parsimonious rule-based model for sentiment analysis of social media text. In: Web and Social Media (2014)
6. Wills, G., Syarif, I., Prugel-Bennett, A.: SVM parameter optimization using grid search and genetic algorithm to improve classification performance. TELKOMNIKA **14**(4), 1502–1509 (2016)
7. Losada, D.E., Crestani, F.: A test collection for research on depression and language use. In: Fuhr, N., Quaresma, P., Gonçalves, T., Larsen, B., Balog, K., Macdonald, C., Cappellato, L., Ferro, N. (eds.) CLEF 2016. LNCS, vol. 9822, pp. 28–39. Springer, Cham (2016). doi:10.1007/978-3-319-44564-9_3
8. Nadeem, M.: Identifying depression on Twitter. arXiv preprint:1607.07384 (2016)
9. Park, M., Cha, C., Cha, M.: Depressive moods of users portrayed in Twitter. In: Proceedings of the ACM SIGKDD Workshop on Healthcare Informatics, pp. 1–8 (2012)
10. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E.: Scikit-learn: machine learning in Python. J. Mach. Learn. Res. **12**, 2825–2830 (2011)
11. Wang, X., Zhang, C., Ji, Y., Sun, L., Wu, L., Bao, Z.: A depression detection model based on sentiment analysis in micro-blog social network. In: Li, J., Cao, L., Wang, C., Tan, K.C., Liu, B., Pei, J., Tseng, V.S. (eds.) PAKDD 2013. LNCS, vol. 7867, pp. 201–213. Springer, Heidelberg (2013). doi:10.1007/978-3-642-40319-4_18
12. WHO. http://www.euro.who.int. Accessed 06 May 2017
13. WHO. Preventing Suicide: A Global Imperative (2015)
14. Srinivasan, P., Yang, C.: Life satisfaction and the pursuit of happiness on Twitter. PLoS ONE **11**(3), e0150881 (2016)

Author Index