

Witold Abramowicz
Adrian Paschke (Eds.)

LNBIP 320

Business Information Systems

21st International Conference, BIS 2018
Berlin, Germany, July 18–20, 2018
Proceedings

bis
BUSINESS INFORMATION SYSTEMS 2018

 Springer

Lecture Notes in Business Information Processing

320

Series Editors

Wil M. P. van der Aalst

RWTH Aachen University, Aachen, Germany

John Mylopoulos

University of Trento, Trento, Italy

Michael Rosemann

Queensland University of Technology, Brisbane, QLD, Australia

Michael J. Shaw

University of Illinois, Urbana-Champaign, IL, USA

Clemens Szyperski

Microsoft Research, Redmond, WA, USA

More information about this series at <http://www.springer.com/series/7911>

Witold Abramowicz · Adrian Paschke (Eds.)

Business Information Systems

21st International Conference, BIS 2018
Berlin, Germany, July 18–20, 2018
Proceedings

Editors

Witold Abramowicz
Poznan University of Economics
and Business
Poznan
Poland

Adrian Paschke
Fraunhofer FOKUS
Berlin
Germany

ISSN 1865-1348 ISSN 1865-1356 (electronic)
Lecture Notes in Business Information Processing
ISBN 978-3-319-93930-8 ISBN 978-3-319-93931-5 (eBook)
<https://doi.org/10.1007/978-3-319-93931-5>

Library of Congress Control Number: 2018947348

© Springer International Publishing AG, part of Springer Nature 2018

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors, and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, express or implied, with respect to the material contained herein or for any errors or omissions that may have been made. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Printed on acid-free paper

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

Preface

During the 21 years of the International Conference on Business Information Systems, the conference has grown to be a well-renowned event for the scientific and business communities. Every year the conference gathers international researchers for scientific discussions on modelling, development, implementation, and application of business information systems based on innovative ideas and computational intelligence methods. The 21st edition of the BIS conference was jointly organised by the Fraunhofer Institute for Open Communication Systems, Germany and Poznań University of Economics and Business, Department of Information Systems, Poland, and was held in Berlin, Germany.

Digital technologies transform the way how business is made, companies grow, links between people are created and evolve, and many others. Digital Transformation means that digital usages inherently enable new types of innovation and creativity. In general, Digital Transformation is described as “the total and overall societal effect of digitalization”. Digitalization resulted in the transformation of existing business models, socio-economic structures, legal and policy measures, organizational patterns, cultural barriers, etc. There are a number of reasons why businesses undergo Digital Transformation. The main argument is that they simply have to.

The BIS conference follows popular research trends, both in the academic and business domain. Thus, to continue this tradition, the theme of BIS 2018 was “Digital Transformation an imperative in today’s business markets” and the goal was to enable sharing theoretical and practical knowledge of the ongoing Digital Transformation activities and induce further innovations that would affect both individual businesses and whole domains, such as administration, communication, art, medicine, healthcare, finance, and science.

The first part of the BIS 2018 proceedings is dedicated to Big and Smart Data and Artificial Intelligence. This is followed by other research directions that were discussed during the conference, including Business and Enterprise Modelling, ICT Project Management, Process Management and Smart Infrastructures. Finally, the proceedings end with Social Media and Web-based Business Information Systems as well as Applications, Evaluations and Experiences of the newest research trends in various domains.

The Program Committee of BIS 2018 consisted of 77 members who carefully evaluated all the submitted papers. Based on their extensive reviews, 30 papers were selected.

We would like to thank everyone who helped to build an active community around the BIS conference. First of all, we want to express our appreciation to the reviewers for taking the time and effort to provide insightful comments. We wish to thank all the keynote speakers who delivered enlightening and interesting speeches. Last but not least, we would like to thank all the authors who submitted their papers as well as all the participants of BIS 2018.

Organization

BIS 2018 was organized by the Fraunhofer Institute for Open Communication Systems and Poznań University of Economics and Business, Department of Information Systems.

Program Committee

Witold Abramowicz (Co-chair)	Poznań University of Economics and Business, Poland
Adrian Paschke (Co-chair)	Fraunhofer FOKUS and Freie Universität Berlin, Germany
Frederik Ahlemann	University of Duisburg-Essen, Germany
Rainer Alt	Leipzig University, Germany
Dimitris Apostolou	University of Piraeus, Greece
Timothy Arndt	Cleveland State University, USA
Sören Auer	TIB Leibniz Information Center Science and Technology and University of Hannover, Germany
Eduard Babkin	INSA Rouen; State University, Higher School of Economics (Nizhny Novgorod), Russia
Morad Benyoucef	University of Ottawa, Canada
Tiziana Catarci	Università di Roma la Sapienza, Italy
François Charoy	Université de Lorraine, LORIA, Inria, France
Rafael Corchuelo	University of Seville, Spain
Christophe Debryne	University College Odisee, Belgium
Josep Domingo-Ferrer	Universitat Rovira i Virgili, Spain
Suzanne Embury	The University of Manchester, UK
Vadim Ermolayev	Zaporozhye National University, Ukraine
Werner Esswein	Technische Universität Dresden, Germany
Anna Fensel	Semantic Technology Institute (STI) Innsbruck, University of Innsbruck, Austria
Agata Filipowska	Poznań University of Economics and Business, Poland
Adrian Florea	Lucian Blaga University of Sibiu, Romania
Johann-Christoph Freytag	Humboldt Universität zu Berlin, Germany
Naoki Fukuta	Shizuoka University, Japan
Jaap Gordijn	Vrije Universiteit Amsterdam, The Netherlands
Volker Gruhn	Universität Duisburg-Essen, Germany
Francesco Guerra	UniMo, Italy
Hele-Mai Haav	Institute of Cybernetics at Tallinn University of Technology, Estonia
Martin Hepp	Universität der Bundeswehr München, Germany

Constantin Houy	Institute for Information Systems at DFKI (IWi), Germany
Christian Huemer	Vienna University of Technology, Austria
Björn Johansson	Lund University, Sweden
Monika Kaczmarek	University of Duisburg Essen, Germany
Pawel Kalczynski	California State University, Fullerton, USA
Kalinka Kaloyanova	University of Sofia, Bulgaria
Naouel Karam	Freie Universität Berlin, Germany
Uzay Kaymak	Eindhoven University of Technology, The Netherlands
Marite Kirikova	Riga Technical University, Latvia
Gary Klein	University of Colorado Boulder, USA
Mathias Klier	University of Ulm, Germany
Ralf Klischewski	German University in Cairo, Egypt
Ralf Knackstedt	University of Hildesheim, Germany
Andrzej Kobylinski	Warsaw School of Economics, Poland
Ryszard Kowalczyk	Swinburne University of Technology, Australia
Marek Kowalkiewicz	Queensland University of Technology, Australia
Eva Kühn	Vienna University of Technology, Austria
Andre Ludwig	Kühne Logistics University, Germany
Leszek Maciaszek	Wrocław University of Economics, Poland
Raimundas Matulevicius	University of Tartu, Estonia
Heinrich C. Mayr	Alpen-Adria-Universität Klagenfurt, Austria
Massimo Mecella	Sapienza University of Rome, Italy
Andreas Oberweis	Karlsruhe Institute of Technology, Germany
Eric Paquet	National Research Council, Canada
Jaroslav Pokorný	Charles University in Prague, Czech Republic
Birgit Proell	FAW, Johannes Kepler University Linz, Austria
Elke Pulvermueller	Institute of Computer Science, University of Osnabrück, Germany
António Rito Silva	Universidade de Lisboa, Portugal
Virgilijus Sakalauskas	Vilnius University, Lithuania
Sherif Sakr	The University of New South Wales, Australia
Demetrios Sampson	Curtin University, Australia
Juergen Sauer	University of Oldenburg, Germany
Stefan Schulte	Vienna University of Technology, Austria
Elmar Sinz	University of Bamberg, Germany
Alexander Smirnov	SPIIRAS, Russia
Stefan Smolnik	University of Hagen, Germany
Andrzej Sobczak	Warsaw School of Economics, Poland
Henk Sol	University of Groningen, The Netherlands
Srinath Srinivasa	International Institute of Information Technology, Bangalore, India
Steffen Staab	Institute WeST, University Koblenz-Landau, Germany and WAIS, University of Southampton, UK
York Sure-Vetter	Karlsruhe Institute of Technology, Germany
Jerzy Surma	Warsaw School of Economics, Poland

Kia Teymourian	Boston University, USA
Genny Tortora	University of Salerno, Italy
Nils Urbach	University of Bayreuth, Germany
Herve Verjus	Université de Savoie, LISTIC, Polytech'Savoie, France
Herna Viktor	University of Ottawa, Canada
Krzysztof Wecel	Poznań University of Economics and Business, Poland
Hans Weigand	Tilburg University, The Netherlands
Mathias Weske	HPI, University of Potsdam, Germany
Anna Wingkvist	Linnaeus University, Sweden
Guido Wirtz	University of Bamberg, Germany

Organizing Committee

Adrian Paschke (Co-chair)	Fraunhofer FOKUS and Freie Universität Berlin, Germany
Bartosz Perkowski (Co-chair)	Poznań University of Economics and Business, Poland
Barbara Gołębiewska	Poznań University of Economics and Business, Poland
Marko Harasic	Fraunhofer FOKUS and Freie Universität Berlin, Germany
Włodzimierz Lewoniewski	Poznań University of Economics and Business, Poland
Milena Stróżyna	Poznań University of Economics and Business, Poland

Additional Reviewers

Anglès-Tafalla, Carles	Kaczmarek, Stefanie
Awad, Ahmed	Kosmol, Linda
Bader, Sebastian	Laifa, Meriem
Bazhenova, Ekaterina	Liutvinavicius, Marius
Braun, Richard	Malyzhenkov, Pavel
Burwitz, Martin	Mogadala, Aditya
Dadashnia, Sharam	Morariu, Daniel
Dittes, Sven	Mulero Vellido, Rafael
Ebner, Katharina	Nikaj, Adriatik
El Shawi, Radwa	Planer, Martin
Ferrarelli, Paola	Pufahl, Luise
Graef, Roland	Radschek, Sophie Therese
Grieger, Marcus	Rehse, Jana
Gutermuth, Oliver	Ribes-González, Jordi
Hassan, Fadi	Ricci, Sara
Hewelt, Marcin	Richter, Peggy
Hornung, Olivia	Schweizer, André
Jentsch, Christian	Sejdovic, Suad
Johannsen, Florian	Thimm, Matthias
Joskowicz, Geri	Wehner, Benjamin
Jöhnk, Jan	Weller, Tobias

Contents

Big and Smart Data and Artificial Intelligence

A Hybrid Approach to Implement Data Driven Optimization into Production Environments	3
<i>Rachaa Ghabri, Pascal Hirmer, and Bernhard Mitschang</i>	
Human Perception of Enriched Topic Models	15
<i>Wojciech Lukaszewicz, Alexandru Todor, and Adrian Paschke</i>	
Predictive Quality: Towards a New Understanding of Quality Assurance Using Machine Learning Tools	30
<i>Oliver Nalbach, Christian Linn, Maximilian Derouet, and Dirk Werth</i>	

Business and Enterprise Modelling

Application of Inductive Reference Modeling Approaches to Enterprise Architecture Models	45
<i>Felix Timm, Katharina Klohs, and Kurt Sandkuhl</i>	
Towards a Typology of Approaches for Sustainability-Oriented Business Model Evaluation	58
<i>Thorsten Schoormann, Anna Kaufhold, Dennis Behrens, and Ralf Knackstedt</i>	
Towards Agility in IT Governance Frameworks	71
<i>Sulejman Vejseli and Alexander Rossmann</i>	
Organizations in Transformation: Agility as Consequence or Prerequisite of Digitization?	86
<i>Dominic Lindner and Christian Leyh</i>	
Information Security Management Systems - A Maturity Model Based on ISO/IEC 27001	102
<i>Diogo Proença and José Borbinha</i>	
Repairing Outlier Behaviour in Event Logs	115
<i>Mohammadreza Fani Sani, Sebastiaan J. van Zelst, and Wil M. P. van der Aalst</i>	

ICT Project Management

Big Data Enabled Organizational Transformation: The Effect of Inertia
in Adoption and Diffusion 135
Patrick Mikalef, Rogier van de Wetering, and John Krogstie

Amalgamation of 3D Printing Technology and the Digitalized
Industry – Development and Evaluation of an Open Innovation
Business Process Model. 148
Danielle Warnecke, Gor Davidovic Gevorkjan, and Frank Teuteberg

Process Management

Fast Incremental Conformance Analysis for Interactive Process Discovery . . . 163
*P. M. Dixit, J. C. A. M. Buijs, H. M. W. Verbeek,
and W. M. P. van der Aalst*

Business Process Compliance and Business Process Change: An Approach
to Analyze the Interactions. 176
Tobias Seyffarth, Stephan Kuehnel, and Stefan Sackmann

Mining Hybrid Business Process Models: A Quest for Better Precision 190
*Dennis M. M. Schunselaar, Tijs Slaats, Fabrizio M. Maggi,
Hajo A. Reijers, and Wil M. P. van der Aalst*

Extending BPSim Based on Workflow Resource Patterns. 206
Nehal Afifi, Ahmed Awad, and Hisham M. Abdelsalam

Towards Implementing REST-Enabled Business Process Choreographies 223
Adriatik Nikaj, Marcin Hewelt, and Mathias Weske

Disambiguation of DMN Decision Tables 236
Kimon Batoulis and Mathias Weske

Smart Infrastructures

Using Blockchain Technology for Business Processes
in Purchasing – Concept and Case Study-Based Evidence 253
Stefan Tönnessen and Frank Teuteberg

Developing a Multiple-Objective Demand Response Algorithm
for the Residential Context. 265
Dennis Behrens, Thorsten Schoormann, and Ralf Knackstedt

Toward Resilient Mobile Integration Processes 278
Daniel Ritter and Manuel Holzleitner

Social Media and Web-Based Business Information Systems

Tight and Loose Coupling in Evolving Platform Ecosystems: The Cases of Airbnb and Uber. 295
Andreas Hein, Markus Böhm, and Helmut Krčmar

On Feeding Business Systems with Linked Resources from the Web of Data 307
Andrea Cimmino and Rafael Corchuelo

Increasing the Explanatory Power of Investor Sentiment Analysis for Commodities in Online Media 321
Achim Klein, Martin Riekert, Lyubomir Kirilov, and Joerg Leukel

Comparative Analysis of the Informativeness and Encyclopedic Style of the Popular Web Information Sources 333
Nina Khairova, Włodzimierz Lewoniewski, Krzysztof Węcel, Mamyrbayev Orken, and Mukhsina Kuralai

Applications, Evaluations, and Experiences

Satellite Imagery Analysis for Operational Damage Assessment in Emergency Situations. 347
German Novikov, Alexey Trekin, Georgy Potapov, Vladimir Ignatiev, and Evgeny Burnaev

Qualitative Assessment of Machine Learning Techniques in the Context of Fault Diagnostics 359
Thilo Habrich, Carolin Wagner, and Bernd Hellingrath

A Comparative Evaluation of Log-Based Process Performance Analysis Techniques 371
Fredrik Milani and Fabrizio M. Maggi

Blockchain for Business Applications: A Systematic Literature Review 384
Ioannis Konstantinidis, Georgios Siaminos, Christos Timplalexis, Panagiotis Zervas, Vassilios Peristeras, and Stefan Decker

ICT-Based Support for the Collaboration of Formal and Informal Caregivers – A User-Centered Design Study. 400
Madeleine Renyi, Frank Teuteberg, and Christophe Kunze

Identifying Suitable Representation Techniques for the Prioritization of Requirements and Their Interdependencies for Multiple Software Product Lines 412
Stephanie Lewellen and Markus Helfert

Author Index 425

Big and Smart Data and Artificial Intelligence



A Hybrid Approach to Implement Data Driven Optimization into Production Environments

Rachaa Ghabri^(✉), Pascal Hirmer, and Bernhard Mitschang

Institute of Parallel and Distributed Systems, University of Stuttgart,
Universitätsstr. 38, 70569 Stuttgart, Germany
{rachaa.ghabri,pascal.hirmer,bernhard.mitschang}@ipvs.uni-stuttgart.de

Abstract. The potential of data analytics to improve business processes is commonly recognized. Despite the general enthusiasm, the implementation of data-driven methods in production environments remains low. Although established models, such as CRISP-DM, offer a structured process in order to deploy data analytics in the industry, manufacturing companies still need to choose a starting point, assess the business benefit, and determine a pragmatic course of action. In this paper, we introduce an approach to handle these issues based on a case study from automotive manufacturing. The results are discussed based on a set of requirements derived from the case study.

Keywords: Data driven optimization · Production environment
Top-down · Bottom-up

1 Introduction

Globalization, shorter product life cycles, and rapidly changing customer needs lead to increasing competitive pressure in the manufacturing industry [1]. In addition to the high product quality and variety, flexibility and short delivery times are also important success factors [2]. Thus, efficient and continuously improved production processes are key prerequisites for a manufacturing company to become and remain successful in the market [3]. In other business sectors, especially e-commerce and internet based services, big data computing and analytics are successfully used for data driven process optimization [4]. This fact puts forth the potential of data driven optimization as a means to boost business processes. Manufacturing companies can also exploit this potential and use data driven optimization in order to meet the ever increasing demands.

In most cases, companies use the *bottom-up approach*, where business-relevant knowledge is searched in all the available data, for example, by using data mining techniques [5]. This approach is characterized by insufficient focus on specific business objectives and strategies of the respective company, as well as relatively

high investments. Furthermore, the bottom-up approach might be over engineered for the use for production processes, since the semantics and structure of the generated data are usually well known.

Another well-established alternative is the *top-down approach*, where selective gathering and analysis of data are conducted solely based on specific business goals [6]. This approach bears the risk of missing business relevant knowledge and leaving lucrative optimization potential unlevered.

In both cases, having access to important amounts of data and disposing of powerful IT-Tools do not necessarily lead to a successful data driven optimization. In addition, companies must also embrace an efficient course of action to ensure (i) a sufficient focus on the strategic company goals, while (ii) still leveraging every business potential, and (iii) maintaining a reasonable expense-benefit ratio. For the mentioned reasons, the bottom-up and top-down approaches do not deliver optimal results. To antagonize these problems, there needs to be a strategic course of action, which combines the advantages and avoids the disadvantages of both. In this paper, we introduce our idea of a hybrid approach to achieve maximum benefit out of data-driven optimization and support it with a real-life case study from the production of car electronics. An in-depth description of the associated steps in order to apply the hybrid approach as a methodology will be delivered in a subsequent paper.

The remainder of this paper is structured as follows: Sect. 2 describes the background and related work. In Sect. 3, a case scenario is described, which is used to derive requirements for our approach. Section 4 contains the main contribution of our paper: a hybrid approach to implement data driven optimization into production environments. In Sect. 5, we discuss to what extent the hybrid approach meets the requirements of the case scenario. Finally, Sect. 6 concludes the paper and gives an outlook on future work.

2 Background and Related Work

In recent years, procedures have been developed to integrate data-driven optimizations into existing IT and process environments of companies. Usually, these approaches are either classified as bottom-up or top-down. This classification originates from the widely used big data pyramid [7,8], which is depicted in Fig. 1. The bottom layer of the pyramid, the data layer, represents low level data, which can be stored into different, distributed, heterogeneous IT systems or even into so-called data lakes [9,10]. The low level data can then be processed and aggregated, for example, by applying data mining techniques [11], in order to generate information, which is represented through the second layer of the pyramid. This information describes interesting, previously unknown patterns in the data. Interlinking this information and combining it with domain specific expertise leads to the third layer of the pyramid, which represents business relevant knowledge. This knowledge can be used as a basis for further actions in order to reach company goals, for example, by purposefully altering the business processes. A pass through the pyramid in Fig. 1 can be carried out either according to the bottom-up or the top-down approach.




Fig. 1. Top-down vs. bottom-up approach

In the bottom-up approach, raw data produced by heterogeneous distributed systems is used as a foundation to derive knowledge that can, for example, lead to the adaptation of processes to increase their efficiency. A precondition for this approach is a holistic, consistent foundation of data to extract or compute information and, consequently, the desired knowledge. For this purpose, data mining techniques can be used in order to recognize interesting patterns in the data. An approved methodology for this is the Knowledge Discovery in Databases (KDD) process as introduced by Fayyad et al. [12]. The bottom-up approach works well in approaches where the data sources and the goals that should be achieved by data analytics are well-known, e.g., when executing previously modeled data flow pipelines [13] or recognizing situations based on context data [14, 15].

However, once a company chooses to apply the bottom-up approach for data driven optimization, it is confronted with a major issue: there is no warranty that all the efforts lead to good results, and, furthermore, there is no reference, which results should be achieved in order to consider the project successful. Furthermore, the recognized patterns in the data may even be misleading, or may be interpreted in a wrong way, leading to no improvements or even to a worsening of the business processes.

On the other hand, the top-down approach builds on specific company goals, which have been derived from a thorough analysis of the enterprise's business processes and IT systems. Based on the specific goals, suitable data and adequate analysis techniques are purposefully selected. The top-down approach is a target-oriented methodology and is more likely to lead to useful results. However, it comes with the risk of missing important information due to the specific, narrow view on the data. Besides, it can be very difficult to decide which data can be considered as relevant to reach the defined goals.

In summary, the bottom-up and top-down approaches have their respective advantages and shortcomings. With the hybrid approach, we aim for in this paper, we use a combination of both approaches to emphasize their advantages and avoid their disadvantages.

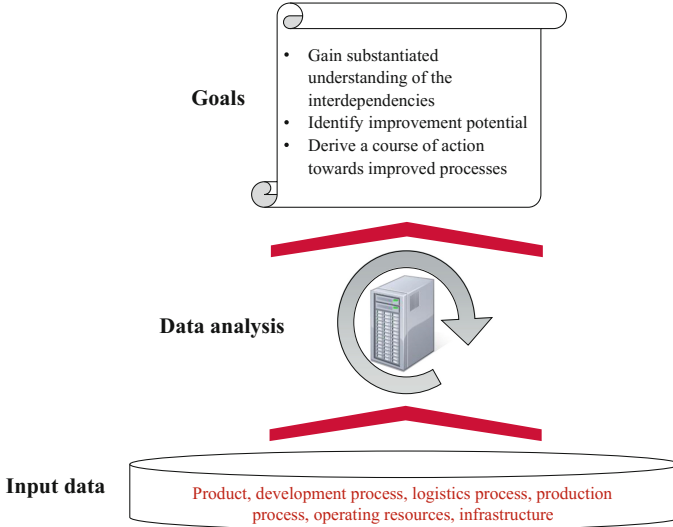


Fig. 2. Using data analytics in the case scenario (bottom-up approach)

3 Case Scenario and Requirements

In this section, we depict the need for a hybrid approach and present one instance of its usage by describing a real-world case scenario from the automotive industry, more specifically from the production of car electronics. Based on this case scenario, we derive a set of requirements for our approach.

3.1 Case Scenario

The increase in complexity of modern cars electronics in terms of architecture, performance, and communication data is one of the reasons why their production processes become more and more challenging for automotive manufacturers [16]. In this case scenario, a large automotive manufacturer aims for improving its electronics production processes in several manufacturing plants. The electronics production processes are the steps of the final assembly, which consist of mounting all the electronic components of the car, flashing the electronic control units with the customer specific software, calibrating the driver assistance systems, commissioning as well as conducting the functional and final testing for all of the electronic components of the car.

In order to improve this part of the product creation process, it is advisable to consider the preceding steps as well. Thus, we are looking at a process chain, which reaches out from technical development, through production planning, up to the operational production. In this case scenario, the process chain is consistently digital, and therefore, generates large amounts of detailed data. This data is usually decentralized, inhomogeneous, and includes detailed information about the product, the development, logistics and production processes,

as well as the used equipment and infrastructure. By the means of data driven optimizations, the automotive manufacturer seeks to gain deep, numerically supported understanding of the interdependencies within the selected business processes and identify improving potentials, as well as deduce an adequate course of action to exploit this potential. Using data analytics (cf. Fig. 2), the manufacturer aims at reaching optimization goals, e.g., a course of action towards improved processes.

At the beginning of this project, the responsible employees of the car manufacturer are confronted with the task of gaining and keeping a perspective of the large amounts of inhomogeneous, apparently incoherent data. Furthermore, several factors make it difficult to decide, which data sets should be considered significant for the analysis. First, the examined business processes show a lack of transparency due to their high complexity and the numerous interdepartmental interfaces. Secondly, the documentation of the processes and the meta data might show some incoherency. Lastly, the conjunction of data sources is not always conducive, which can undermine data with a potential for business relevant knowledge. These conditions render the data preprocessing and integration a cumbersome task that can also affect the motivation.

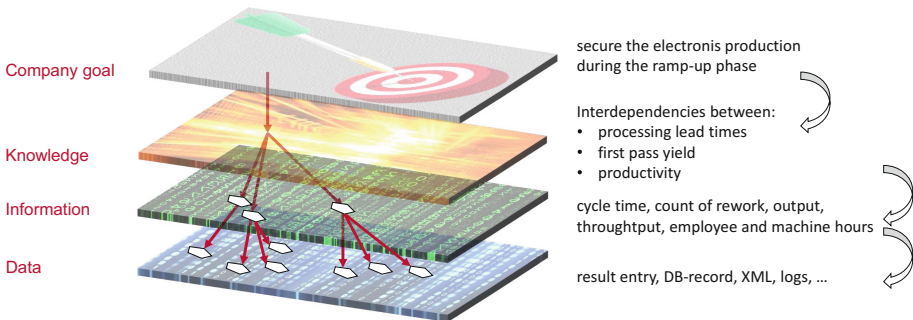


Fig. 3. Top-down approach for the case scenario

As mentioned in Sect. 2, the bottom-up approach starts with collecting and storing all available data. The architectural and structural components of an adequate IT-solution for this job are usually a huge financial investment, while it is not clear, whether process improvements would occur and if they would account for company goals with high priority. This uncertainty makes the budget clearance for such projects more difficult. Additionally, the company would have to commit to a certain IT-solution prior to conducting a spike test to ensure the feasibility and the suitability. Here lies a risk that the IT-solution would emerge as unfit for the business and operating environment of the company [17].

Applying the top-down approach in this case scenario would start with defining a specific company goal, e.g. securing the electronics production during the ramp-up phase of a new car model. Based on this given goal, the necessary

knowledge needs to be determined. For instance, the interdependencies between the processing lead times, the first pass yield, and the productivity of the whole production plant or of a single work station should be useful. The next step is to specify the information, which would lead to the needed knowledge. In this case, the information about the cycle time of a single operation, the count of rework and mistakes, the output, the throughput and the number of employee and machine-hours should be considered. Up to this point of the data driven optimization, it is irrelevant, which IT-tools and architectural components will be used. For the transition from the information layer to the data layer it is, however, necessary to determine the right data sources, design the information model including metadata management, as well as to conduct the adequate data processing. Therefore, it is necessary to make a decision about the IT-solution to be employed. Using the top-down approach for this example is depicted in Fig. 3.

While the top-down approach is more likely to help the company reach its optimization goals, its scope of action is limited to one single issue. Thus, using this approach only allows a parochial view of the data and the improvement potentials, rather than considering the wider context. For instance, when the top-down approach is used in the example above, important insights, such as the impact of the infrastructure on the processing lead time, would remain undetected. Besides, by focusing on a given goal, the domain experts in the company miss out on an opportunity to expand their domain knowledge and discover previously unknown interdependencies within the process chain.

In summary, both approaches cannot provide a satisfying solution for the usage of data analytics in order to improve production processes of this or other scenarios. In the following section, we derive a set of requirements to cope with the mentioned issues, which build the foundation of our approach.

3.2 Requirements

The hybrid approach, we aim for in this paper, minimizes the risks and combines the advantages of the approaches described above. We define the following requirements for our approach:

- (R1) Contribution to high-priority company goals: The hybrid approach needs to ensure that the data driven optimization is set up to contribute to strategic, highly prioritized goals of the company. Thus, the first step of the approach must consist in defining a concrete outcome of the project. By doing so, it is possible to evaluate and rank a specific data driven project based on the company's current priorities.
- (R2) Full development of the potential for improvement: The hybrid approach must ensure that the data analysis reveals every worthy room for improvement: as a counterpart for the pragmatic implementation, the long-term expectation out of data driven optimization is to look into every potentially value-adding insight.

- (R3) Optimal cost-benefit ratio: The hybrid approach aims to achieve an optimal cost-benefit ratio out of data driven optimization: it avoids investments with a long payback period. Instead, it relies on incremental investments with many “low-hanging fruits”.
- (R4) Promotion of feasibility: The approach must promote the feasibility of data driven optimization within the business- and operating-guidelines of the company: data driven optimization is not conducted for its own sake, but rather to bring a practical benefit for the company. Therefore, they should be conducted pragmatically and with minimum distortion of the core business.

4 Hybrid Approach for Data Driven Optimization

The goal of this paper is an approach to implement data driven optimization into production environments, while minimizing the disadvantages and highlighting the advantages of the established bottom-up and top-down approaches. The hybrid approach consists of a purposeful, structured alteration and combination of the top-down and bottom-up approaches in order to join a motivating effectiveness with a holistic performance, and at the same time to avoid high, uncertain investment. Initially, a set of use cases are concluded in the style of the top-down approach. In addition to fulfilling the specific purpose of the use case, each successful execution will reveal business-valuable data sets. These are the data sets which evidentially lead to profitable knowledge for the company. Such data sets are referred to as *data treasures* in the context of this paper. Using the bottom-up approach, the data treasures are then analyzed and the contained information is correlated in order to gain insights beyond the discrete use cases. In doing so, companies can ensure a maximum benefit out of data-driven optimizations while holding the risks at a viable level. Figure 4 shows the steps of the hybrid method, which is explained in the following sections.

4.1 Derive, Prioritize and Execute Use Cases

The first phase of the hybrid approach is based on the top-down approach. As mentioned in Sect. 2, this approach begins with the definition of a business goal that is in line with the company’s strategic objectives. Queries with a direct reference to the production field are then derived from the business goal. The queries should be formulated as precisely as possible, and the corresponding frame conditions should be specified in order to answer them in the context of specific use cases by using analytics techniques. We suggest the following references in order to convert a query into a manageable use case:

- Which type of analytics, i.e., descriptive, diagnostic, predictive, or prescriptive [18], is suitable to answer the respective query?
- What are the key performance indicators, parameters and influencing factors involved in the query?
- Is the underlying data already available and, if not, what needs to be done to make it available?

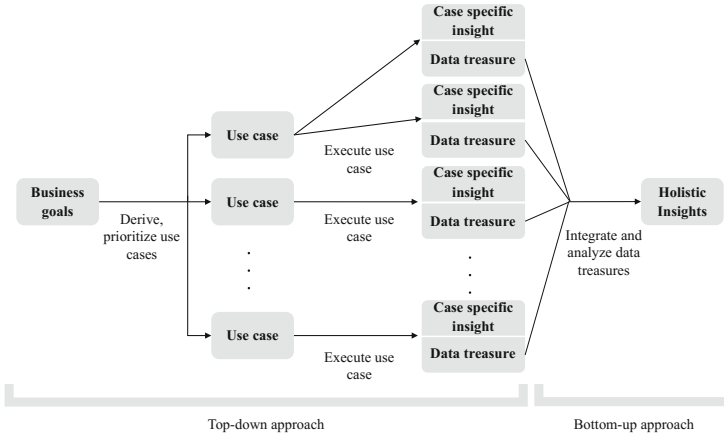


Fig. 4. Steps of the hybrid approach

- Which are the sources of the underlying data and which format does the raw data have?
- Which requirements must the data processing meet, e.g., real-time or incremental processing?
- How long is the period of time that is considered in the analysis and how frequently will the analysis be conducted?
- What practical benefit for the company comes with answering the query?

After converting a query into a use case and based on the answers to the questions above, a potential analysis is to be conducted. At this point, we recommend to look into the following features to assess the priority of a given use case: (i) acuteness, defining to what extent the use case attends to urgent issues of the company, (ii) feasibility, describing how much effort goes into providing and processing the needed data, and (iii) relevance, examining the bearing of the benefit. The potential analysis helps the company identify result-oriented, data-based use cases in an efficient, structured and repeatable manner.

For the highly-prioritized use cases, the required data sources are made available, access authorization is managed, and data security measures are taken. Afterwards, the data is processed, e.g., through validation, cleaning, and aggregation, in order to prepare it for the subsequent analysis. In the analysis step, statistical evaluation is used in order to answer the query with the help of the data. The results of the analysis are then made comprehensible by means of appropriate visualization. The latter is then evaluated by the domain experts and used as support to conclude a course of actions. The sequence of the first phase of the hybrid approach is depicted in Fig. 5.

The first phase of the hybrid approach are conducted in a cyclical manner. This means that the output of a successfully executed use case, i.e., the gained in-sights, may influence the input of the next use case, for example, through the adaptation of a defined business goal or the definition of new ones. The




Fig. 5. The first phase of the hybrid approach (Steps in red color are conducted by the business department, steps in grey color by the IT department) (Color figure online)

execution of the first phase of the hybrid approach calls for the collaboration of the respective business departments and the IT department. In Fig. 5, the steps marked in red color are to be conducted by the business department, while the steps marked in grey color are the tasks of the IT department. To reach maximum benefit, it is recommended to comply with the allocation of tasks, so that each department can concentrate on its core expertise.

4.2 Integrate and Analyze Data Treasures

The executed use cases conduce to identify the parts of the data jungle, which contain information with business value. As already mentioned, this data shall be referred to as data treasure. Once a data treasure has been identified, it is made available in a central data storage, for example, a data lake, (cf. Fig. 6). This way, the central data storage, i.e. the data lake, will only contain data with confirmed usefulness, and will expand with every conducted use case. The data treasures of a specific use case, which show mutual correlation, are then assigned to one cluster and should be considered as a coherent entity.

4.3 Collecting and Accessing Data Treasures

In the style of the bottom-up approach, the established entities are then examined in order to find correlations with each other or with further parameters from different yet related use cases. By doing so, the analysis is carried out not only within the boundaries of single use cases, but rather on a holistic level. Since

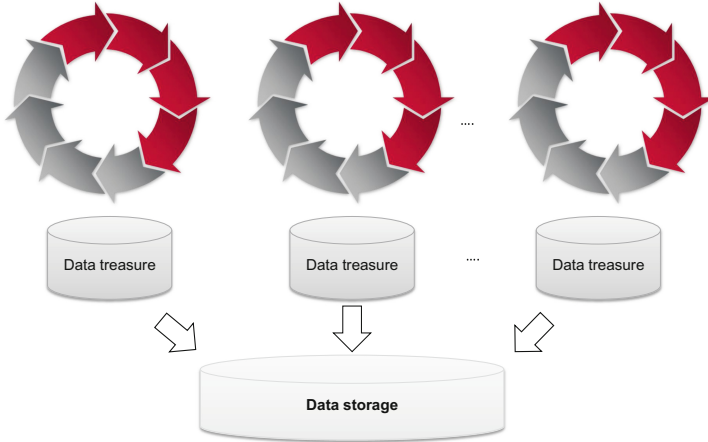


Fig. 6. Collecting and accessing data treasures

this step is likely to be sophisticated and costly, it should be ensured that the efforts are well-invested. For that reason, the risk of trailing away with irrelevant or pseudo-correlations needs to be minimized. We recommend this step of the hybrid approach to adhere specific frame conditions in order to maintain the efficiency. For instance, integrating and analyzing the data treasures can be carried out for a specific period of time, a specific car model, or a specific manufacturing technology.

5 Discussion

In this section, we discuss our approach in terms of fulfilling the requirements from Sect. 3.2. One of the main features of the hybrid approach is its pronounced goal-orientation. The first phase of the hybrid approach ensures that careful consideration is dedicated to defining and selecting project objectives that are in line with the company's goals in order to stay focused on what is strategically important; hence, the first requirement (R1), i.e. contribution to high-priority company goals, is fulfilled. Nevertheless, the hybrid approach allows for exploiting the potential of data analysis beyond rigidly set objectives. Companies can reach high-level value through the purposeful application of the bottom-up approach, which makes sure that none of the potentials for data-driven optimization remain undiscovered. Therefore the hybrid approach also meets the requirement of allowing a full development of the potential for improvement (R2). Besides, as the phases of the hybrid approach are meant to be executed consecutively, the company will have the possibility to gradually ascertain the true business value of the available data sources; to concentrate on utilizing data analysis as a means to improve the business processes. Since the hybrid approach initially relies on the consecutive implementation of several stand-alone use cases, it does

not call for a primary large investment. It rather favors gradual investments with perceptible impact. Furthermore, the company is able to avoid committing to a costly, sophisticated IT-solutions before thoroughly investigating the specific circumstances. For these reasons the hybrid approach is in line with the requirement (R3) of achieving an optimal cost-benefit ratio. In terms of the requirement of promoting the feasibility (R4), the hybrid approach is characterized by the sensible, practical usage of data analysis in production environments. Due to the sequence of its phases, the hybrid approach provides the company with the opportunity to readjust its course of action in the manner of a loop control-system. Moreover, the design of the hybrid approach allows to achieve quick-wins, which accounts for a sense of achievement among the involved employees and results in a higher motivation.

6 Conclusion and Future Work

Data driven optimization is an effective, innovative method for revealing interdependencies and detecting anomalies within the production processes, in order to make them more transparent, stable and controllable. However, a pragmatic, goal oriented and yet holistic approach is key to deploy the full potential of this method. This can be accomplished by adapting and combining the top-down and bottom-up approach. In this paper, we explained the potential analysis of data driven optimization in the production environment and introduced our idea for a hybrid approach for implementing it. In future work, we will deliver further details of the concepts as well as introduce an in-depth method to apply them. For instance, we will depict applicable approaches integration approaches and techniques, in order to interlink the case-specific data-treasures. Furthermore, we intend to look into quantifying the advantage of the hybrid approach in comparison with the conventional top-down and bottom-up approaches in terms of explicit figures.

References

1. Khan, A., Turowski, K.: A survey of current challenges in manufacturing industry and preparation for industry 4.0. In: Abraham, A., Kovalev, S., Tarassov, V., Snášel, V. (eds.) Proceedings of the First International Scientific Conference “Intelligent Information Technologies for Industry” (IITI 2016). AISC, vol. 450, pp. 15–26. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-33609-1_2
2. Jacob, F., Strube, G.: Why go global? The multinational imperative. In: Abele, E., Meyer, T., Näher, U., Strube, G., Sykes, R. (eds.) Global Production, pp. 2–33. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-71653-2_1
3. Slack, N., Chambers, S., Johnston, R.: Operations Management. Pearson Education, New York (2010)
4. Bryant, R., Katz, R.H., Lazowska, E.D.: Big-data computing: creating revolutionary breakthroughs in commerce, science and society (2008)
5. Raval, K.M.: Data mining techniques. Int. J. Adv. Res. Comput. Sci. Softw. Eng. **2**(10) (2012)

6. Couldry, N., Powell, A.: Big data from the bottom up. *Big Data Soc.* **1**(2), 2053951714539277 (2014)
7. Lovelace, R.: *The data revolution: big data, open data, data infrastructures and their consequences*, by rob kitchin. 2014. Thousand Oaks, California: Sage Publications. 222+XVII. ISBN: 978-1446287484. *J. Reg. Sci.* **56**(4), 722–723 (2016)
8. McAfee, A., Brynjolfsson, E., Davenport, T.H., Patil, D., Barton, D.: Big data: the management revolution. *Harvard Bus. Rev.* **90**(10), 61–67 (2012)
9. Gartner: Data lake. Gartner IT Glossary (2017)
10. Held, J.: Will data lakes turn into data swamps or data reservoirs? (2014)
11. Han, J., Kamber, M., Pei, J.: Mining frequent patterns, associations, and correlations. In: *Data Mining: Concepts and Techniques*, 2nd edn., pp. 227–283. Morgan Kaufmann Publishers, San Francisco (2006)
12. Fayyad, U., Piatetsky-Shapiro, G., Smyth, P.: The KDD process for extracting useful knowledge from volumes of data. *Commun. ACM* **39**(11), 27–34 (1996)
13. Hirmer, P., Behringer, M.: FlexMash 2.0 - flexible modeling and execution of data Mashups. In: Daniel, F., Gaedke, M. (eds.) *RMC 2016*. CCIS, vol. 696, pp. 10–29. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-53174-8_2
14. Hirmer, P., Wieland, M., Schwarz, H., Mitschang, B., Breitenbücher, U., Sáez, S.G., Leymann, F.: Situation recognition and handling based on executing situation templates and situation-aware workflows. *Computing* **99**, 163–181 (2017)
15. Wieland, M., Hirmer, P., Steimle, F., Gröger, C., Mitschang, B., Rehder, E., Lucke, D., Rahman, O.A., Bauernhansl, T.: Towards a rule-based manufacturing integration assistant. In: Westkämper, E., Bauernhansl, T. (eds.) *Proceedings of the 49th CIRP Conference on Manufacturing Systems (CIRP-CMS 2016)*, Stuttgart, Germany, 25–27 May 2016, *Procedia CIRP*, vol. 57, pp. 213–218. Elsevier, January 2017
16. Hermann, M., Pentek, T., Otto, B.: Design principles for industrie 4.0 scenarios. In: *2016 49th Hawaii International Conference on System Sciences (HICSS)*, pp. 3928–3937, January 2016
17. Tönne, A.: Big Data is no longer equivalent to Hadoop in the industry. In: *Proceedings of 17. Datenbanksysteme für Business, Technologie und Web (BTW) (2017)*
18. Hagerty, J.: *2017 planning guide for data and analytics* (2016)



Human Perception of Enriched Topic Models

Wojciech Lukaszewicz, Alexandru Todor^(✉), and Adrian Paschke

AG Corporate Semantic Web, Institute for Computer Science,
Freie Universität Berlin, 14195 Berlin, Germany
{wojlukas,todor,paschke}@inf.fu-berlin.de

Abstract. Topic modeling algorithms, such as LDA, find *topics*, hidden structures, in document corpora in an unsupervised manner. Traditionally, applications of topic modeling over textual data use the bag-of-words model, i.e. only consider words in the documents. In our previous work we developed a framework for mining enriched topic models. We proposed a bag-of-features approach, where a document consists not only of words but also of linked named entities and their related information, such as types or categories.

In this work we focused on the feature engineering and selection aspects of enriched topic modeling and evaluated the results based on two measures for assessing the understandability of estimated topics for humans: model precision and topic log odds. In our 10-model experimental setup with 7 pure resource-, 2 hybrid words/resource- and one word-based model, the traditional bag-of-words models were outperformed by 5 pure resource-based models in both measures. These results show that incorporating background knowledge into topic models makes them more understandable for humans.

1 Introduction

Topic modeling, similarly to text classification, is an established and thoroughly researched field in computer science. Traditionally, both techniques are based on a bag-of-words (BOW) document representation, where one feature corresponds to one word (its stem or lemma), i.e. the word order doesn't count, only the frequencies. As Gabrilovich and Markovitch [3] describe the state of the art in text classification in their 2005 paper, “*after a decade of improvements, the performance of the best document categorization systems became more or less similar, and it appears as though a plateau has been reached [...]*”. For this reason, researchers started working on developing different approaches. Considering the limitations of BOW model, the most natural idea was to enhance the method of document representation.

Scott and Matwin [11] did one of the first efforts of feature vector engineering for the purpose of text classification by using WordNet, a lexical database for english, and converting documents to feature vectors based on this new representation. Recent papers, such as Garla and Brandt [4] and Zong et al. [14],

employ semantic information during the feature engineering step and apply machine learning techniques to learn text categories.

These attempts inspired us to perform feature engineering in the context of topic modeling. We want to incorporate semantic information in order to extend the traditional bag-of-words approach into a novel bag-of-features approach when preparing feature vectors. We plan to consider not only words but also disambiguated Named Entities linked to DBpedia resources and several related entities.

The underlying idea and motivation for our work is based on the fact that topic modeling algorithms draw their information based on the frequencies and co-occurrences of tokens in single documents and across the whole corpus. Because of that, we formulated a hypothesis that, in thematically related documents, the entities and/or their types, hypernyms or categories of corresponding Wikipedia articles should also be overlapping and thus summed frequencies of these terms should be more meaningful and lift up their relevance in discovered topics.

For example, consider a text snippet from a Spiegel Online¹ article that a human would assign a label “politics”: *“Barack Obama is only passing through Germany on his trip to Europe later this week and does not plan to hold substantial talks with Angela Merkel. The White House views the chancellor as difficult and Germany is increasingly being left out of the loop”*. The word *politics* itself has a zero frequency. But if we perform Named Entity Recognition and Disambiguation, the entities *Barack Obama* and *Angela Merkel* will be considered politicians thanks to the enrichment we perform.

In this work we present an approach of mining topic models enriched with background knowledge. We focus on the feature engineering aspect of topic modeling and leave the underlying generative statistical model intact. We assess the quality of this approach based on the evaluation strategy which consists of inspecting the internal coherence of topics and the topic-document assignments in terms of human understanding.

2 Related Work

In contrary to pure word-based LDA algorithm and its variations (such as different sampling techniques or online learning proposed by Hoffman et al. [5] which enables streaming-wise model mining and is thus much less resource-hungry) or applications (Gibbs-sampling-based LDA for gene prediction [10]), topic modeling approaches using enriched feature vectors have not been subject to much research so far.

One of the first methods that contributes to topic modeling using entities instead of words as features has been published by Newman et al. [8]. The authors propose five new models that modify LDA in order to learn pure entity-topic models. They evaluate them with regard to entity prediction and not their information-theoretic value, e.g. by measuring perplexity.

¹ <http://www.spiegel.de/international/>.

Hu et al. [6] present an approach of taxonomy-based topic representation that focuses on entities from a knowledge base. They developed a probabilistic model, Latent Grounded Semantic Analysis, that infers both topics and entities from text corpora and grounds the information to a KB.

Todor et al. [12], our previously published work, approaches enriched topic models in a different way and sees them as predictors for multi-labeled classification tasks. The approach was evaluated on news articles, each of which was labeled with a category. After having mined the topics, we let the model predict the coverage for every document and counted a histogram for every topic of how many times it was most relevant for a particular label (e.g. topic 1 was most relevant for 100 documents about sport, 759 times about politics, etc.). Then, we took the highest value of the label-histogram and from this moment on considered it the label of this topic. For the evaluation we evaluated the classification accuracy. To be more specific, we counted, which (1st, 2nd or 3rd) most relevant topic was the correct one, i.e. associated with the article’s label. The results showed that for every dataset there was at least one enriched topic (consisting of words + linked entities) that outperformed the classic topic consisting solely of words when looking only at the single most relevant predicted topic which was a very positive and important outcome. When we left words aside and only considered linked entities, we had to take also the 2nd and 3rd most relevant topic into account. We explain it with the fact that the vocabulary of linked entities is much smaller compared to words and it is harder to make an unambiguous prediction. On the other hand, feature combinations that have a comparable cumulative accuracy to words within top three predictions, operate on a smaller vocabulary (which has advantages of lower time and space complexity).

We already mentioned *perplexity* (or equivalently *predictive likelihood*) as the established qualitative method for expressing the quality of a topic model. To calculate perplexity on a held-out set D of test documents in LDA, we will characterize the model by the topic matrix Φ and the hyperparameter α , the Dirichlet prior for the document-topic distribution. Hence, what we search for is calculating the log-likelihood of generating every document $d \in D$ given two above-mentioned parameters:

$$LL(D) = \log p(D|\Phi, \alpha) = \sum_{d \in D} \log p(d|\Phi, \alpha)$$

Computed log-likelihood can be used to compare topic models – the higher, the better the model.

For LDA, to take the size of the vocabulary into account, we define perplexity as the exponential of negative log-likelihood divided by the number of tokens (note that the quality of the model increases while perplexity decreases):

$$perplexity(D) = \exp\left(\frac{-LL(D)}{\#tokens}\right)$$

Wallach et al. [13] published an overview of evaluation methods for topic models. They address certain challenges, such as the difficulty of estimating $p(d|\Phi, \alpha)$ and propose using sampling to overcome it.

Another interesting method of evaluating topic models is *coherence*, i.e. examining the existence of a single semantic concept that enfolds the words of the topic. This task can be performed quantitatively and there exist two state-of-the-art methods of calculating coherence – an intrinsic (that doesn’t use external source of information) and an extrinsic one (that might employ external data or statistics to calculate the coherence score). Both methods are based on the same idea of calculating a sum of scores for every pair of top n words for a given topic t :

$$coherence_t = \sum_{i < j} score(w_{t,i}, w_{t,j})$$

The difference between the methods is the score function.

The most popular extrinsic measure is the *UCI* measure, proposed by Newman et al. [9]. The pairwise score function calculates the Pointwise Mutual Information:

$$uci(w_i, w_j) = \log \frac{p(w_i, w_j)}{p(w_i)p(w_j)}$$

The probabilities $p(w)$ and $p(w_i, w_j)$ indicate the probabilities of seeing a word w and a co-occurring pair of words w_i and w_j in a random document. These probabilities are manually estimated using a document corpus different from the one used to mine topics (hence the name, *extrinsic* method), e.g. as the document frequencies of word/word pairs.

The most popular intrinsic coherence measure is *UMass*, posted by Mimno et al. [7]. The proposed score function is a smoothed variant of the conditional log-probability:

$$umass(w_i, w_j) = \log \frac{D(w_i, w_j) + 1}{D(w_i)}$$

The UMass score measures the goodness of a common word being a predictor for a less common one. Here we consider the relations among the top n words from a given topic.

Both measure names come from the Institutes where the authors worked at the times of publications (University of Massachusetts for David Mimno and University of California, Irvine for David Newman).

Let us present another evaluation technique, proposed by Chang et al. [2]. They wanted to qualitatively measure the coherence of the estimated topic models. The state-of-the-art quantitative coherence-based evaluation methods, UMass and UCI, follow a similar high-level idea. A method proposed by Chang et al. abstracts from them and postulates a purely manual evaluation of the *human judgement*, which can be seen as understanding and coherence from the human perspective.

The authors have defined two tasks – *word intrusion* and *topic intrusion*. First one is expected to measure how strongly the most relevant words for the topic compose a coherent semantic concept, an unbreakable unit. For this they draw a random topic from the model, take its five most relevant words and add one top word from remaining topics to create a set of 6 words which they

shuffle and present to a human who is expected to select the intruder. In order to quantitatively measure how well the topics match human understanding, Chang et al. introduced *model precision*, i.e. to which extent the intruders selected by humans correspond to the “real” intruders. Model precision is defined as the fraction of correctly selected intruders (Eq. 1, where w_s is the word selected by the evaluation subject and w is the real intruding word), and thus ranges from 0 (worst) to 1 (best).

$$MP = \sum \mathbb{1}(w_s = w) / S \quad (1)$$

The second task measures how understandable the topics are in terms of assigning them to a given text. To prepare a topic intrusion question the authors draw a random article and consider its topical coverage. They take three most relevant topics and one irrelevant topic, each represented by its eight top words, shuffle them and present to an evaluation subject to select the intruder. The results for this task are evaluated using *topic log odds*. This measure, introduced by Chang et al. measures how good were the guesses of humans. In topic intrusion task, every answer (a topic) has certain probability of generating a given document. Topic log odds sums and normalizes differences of logs of probabilities of real intruder belonging to the document and the intruder selected by the evaluation subject. Intuitively, this way of evaluating makes a lot of sense, since it doesn’t binarily count right/wrong answers, but works as a kind of *error function*.

$$TLO = (\sum \log \hat{\theta}_{d,*} - \log \hat{\theta}_{d,s}) / S \quad (2)$$

Simplifying the notation of Chang et al., Eq. 2 is the definition of topic log odds, where $\hat{\theta}_{d,*}$ is the probability of the intruder topic belonging to the document and $\hat{\theta}_{d,s}$ is the probability of the topic selected by the subject belonging to the document. Because latter is greater or equal to the former, the topic model in terms of *TLO* is better when *TLO* is higher (closer to 0).

3 Approach

Our approach differs in several ways from the state of the art methods using entities in mining topic models. First of all, we do not modify the underlying probabilistic generative model of LDA and can therefore apply our method on any variation and implementation of the algorithm. Second of all, we mine topics that contain named entities linked to a knowledge base and might be used for knowledge acquisition purposes, e.g. taxonomy extraction or knowledge based population. Moreover we only employ one KB – DBpedia and DBpedia Spotlight as the NERD tool and focus on finding the best topic models in this setup. Lastly, we combine two evaluation techniques – we concentrate on achieving low perplexity, as well as measure human perception and interpretability of mined models.

4 Evaluation

We evaluated our approach on three datasets: BBC², The New York Times Annotated Corpus³ (NYT) and DBpedia Abstracts.

The BBC dataset is a collection of 2225 selected news articles from BBC news website⁴ which correspond to stories from five domains (*business, entertainment, politics, sport* and *tech*), published in 2004 and 2005.

NYT is a collection of over 1.8 million selected New York Times articles that span for 20 years, from 1987 to 2007. Over 650 thousand of them have been manually summarized by the library scientists, 1.5 million have been manually tagged with regard to mentions of people, organizations, locations and topic descriptors and over 250 thousand of them have been algorithmically tagged and manually verified. Additionally Java utilities for parsing and processing the corpus which is provided in XML format are included in the download. For these reasons it is one of the most widely used datasets for approaches from the domain of natural language processing. We didn't take the full NYT dataset but reduced it to over 46000 articles which have been pre-categorized into at least one of following ten taxonomy classes: *science, technology, arts, sports, business, health, movies, education, u.s., world*.

We already introduced the small and medium-sized datasets we used for the evaluation. We chose one more dataset that can be categorized as big. Namely it is the corpus of DBpedia abstracts – first paragraphs of Wikipedia articles extracted as free text. Abstracts are connected to the DBpedia resources using the property `abstract` from the DBpedia ontology (<http://dbpedia.org/ontology>). The dataset *Long Abstracts* is available to download from the DBpedia downloads website⁵.

After cleaning the documents and annotating the datasets (some articles contain characters which cause Spotlight to fail) we end up with the numbers of articles per data source displayed in Table 1.

Table 1. Sizes of datasets used in the evaluation.

Dataset	Size
BBC	2224
NYT	46406
Abstracts	4641642

We decided to use 10 feature combinations in our evaluation. It means that every topic model will be estimated 10 times, once for every combination. Let us enumerate them: first of all, we chose to mine models using words

² <http://mlg.ucd.ie/datasets/bbc.html>.

³ <https://catalog.ldc.upenn.edu/ldc2008t19>.

⁴ <http://news.bbc.co.uk/>.

⁵ <http://wiki.dbpedia.org/Downloads2015-10>.

(*w* – an abbreviation in parentheses after every mentioned feature combination will be sometimes used to avoid frequent repeating of long descriptors) to have a possibility to compare models of our enriched approach to a classical one. Second of all, we decided to use words alongside with entities (*we*) and hypernyms (*weh*) – we assumed that linked named entities recognized in text represent an important concept which can be characteristic for the given text. For the same reason we also include hypernyms, which can generalize and “group together” semantically related entities.

The remaining seven models consist purely of DBpedia resources. They include models having entities (*e*), hypernyms (*h*) and both features together (*eh*). The justification for the choice of these three feature sets is analogous – we assume the entities to represent important concepts from a text and hypernyms to generalize them. Next model consists of types only (*t*). We want to see if `rdf:types` of entities recognized in a document are descriptive enough for its content. Last three models contain types alongside with entities, hypernyms and the combinations of both (*et*, *th*, *eth*). Here we wanted to investigate if these three types of features combined account to the quality of mined models.

We decided to omit the models containing categories (`subject` properties from Dublin Core terminology which correspond to Wikipedia categories of a given entity) in the evaluation since we expected them to be too broad and not enough descriptive.

The goal of this work was not finding best numbers of topics through exhaustive evaluation. In this work we want to measure changes in the human perception of enriched topics for different feature sets. However, we needed to determine *k* for every dataset and use it for generating *word/entity* – and *topic intrusion* tasks. For this reason we evaluated perplexity for diverse models (diverse *feature set/number of topics* combinations), as proposed in [1, 13].

For generating the study we wanted to choose sets of models with *k*’s, for which the pure word-based models’ perplexities stop falling. We expected it to allow us to compare models mined using our enriched approach to a *best* model for the given dataset.

We had troubles selecting number of topics for the DBpedia Abstracts dataset since generating models on such a huge corpus is very time consuming (10 h on an 8 core CPU). We were not guaranteed that estimating more and more models with larger numbers of topics would result in finding one where the perplexity would finally start to converge in a reasonable amount of time. Overview of perplexity values for different combinations of *k*’s and feature sets for DBpedia Abstracts corpus can be found in Table 4. We noticed that perplexity for the traditional model mined on words was almost the same for 3000 and 5000 topics.

Eyeballing the questionnaires generated from NYT and DBpedia Abstracts models made us conclude that the improvements in perplexity for models with higher numbers of topics, don’t necessarily carry improvement in the quality of questions.

In the end we decided to use following number of topics per dataset:

- New York Times Annotated Corpus: 125 topics
- BBC Dataset: 30 topics
- DBpedia Abstracts Dataset: 1000 topics

To conduct the survey and manage its results at first we prepared a Google Doc with 60 entity and 60 topic intrusion tasks but we quickly realized it’s much too cumbersome and time consuming to complete – our first subject was not ready after 3h. To address it, we implemented a web service where we were displaying one entity and one topic intrusion task at once instead. Each was randomly chosen from the set of tasks with lowest number of answers so that we keep the answer distribution as balanced as possible. Conducting the evaluation in this manner is acceptable since Chang et al. did not profile their results per user neither (Tables 2 and 3).

Table 2. Perplexity of topic models mined on the BBC dataset using different feature types and numbers of clusters.

	w	w-e	w-e-h	t	t-h	h	e	e-h	e-t	e-t-h
5	212.81	189.44	203.54	18.48	46.51	42.91	26.35	46.26	30.86	52.96
7	209.08	190.04	201.93	16.78	42.03	44.11	26.03	45.54	31.25	49.15
10	201.17	183.82	198.32	15.38	44.41	36.96	24.76	43.87	29.61	42.27
15	195.5	169.47	181.49	14.64	33.05	40.52	23.82	44.01	29.34	45.24
20	186.19	169.71	176.01	11.71	35.38	39.5	23.6	41.67	25.48	44.45
25	182.92	166.79	179.98	11.96	30.98	36.82	23.07	41.23	27.69	38.86
30	180.93	164.36	166.14	11.09	29.36	31.73	22.65	39.73	26.13	37.96
40	181.16	162.23	171.6	10.56	31.29	33.93	22.42	40.29	28.79	35.51

4.1 Results

We sent out the link to the questionnaire among fellow computer science/mathematics students, i.e. people with a bigger than average technical sense and understanding of algorithmic and mathematical concepts. This fact might have been reflected in final results. We closed the survey after receiving 600 answers. It corresponds to 300 answers for each task type and thus 5 answers for each single task. Our study had 10 participants, a similar number compared to 8 in experiments by Chang et al.

We started the evaluation by assessing the model precision and topic log odds overall, as well as for each dataset separately. Out of 300 entity intrusion answers, 260 were correct, which accounts for the model precision of 0.87. We achieved the overall topic log odds of -4.23 . The value itself is not very expressive, but we see that the BBC dataset performs best. All topic log odds and model precision values are presented in Table 5.

Table 3. Perplexity of topic models mined on the New York Times Annotated Corpus dataset using different feature types and numbers of clusters.

	w	w-e	w-e-h	t	t-h	h	e	e-h	e-t	e-t-h
10	296.11	360.62	357.94	8.45	31.31	37.95	188.16	137.78	89.47	99.45
15	281.65	342.87	338.42	8.01	29.64	35.78	167.52	125.37	83.67	90.93
20	272.64	327.23	323.54	7.66	28.1	34.69	157.89	117.67	77.0	85.33
30	260.18	309.92	307.96	7.28	26.43	32.7	140.81	109.04	70.32	79.2
50	247.18	291.02	285.83	7.0	25.04	30.52	124.24	97.1	63.38	70.96
100	232.15	266.39	262.6	6.81	23.72	28.48	109.24	84.63	56.74	63.31
125	226.69	260.96	257.15	6.77	23.41	27.93	104.98	81.31	54.95	61.11
150	223.7	256.49	251.39	6.74	23.13	27.53	101.93	79.28	53.7	59.6
200	219.25	249.08	244.57	6.72	22.81	26.86	98.16	75.74	52.33	57.23
300	213.48	240.5	235.34	6.69	22.45	26.35	92.3	71.94	50.21	54.4
500	207.59	231.28	225.71	6.67	22.05	25.69	87.81	67.13	48.21	51.36

Table 4. Perplexity of topic models mined on the DBpedia Abstracts dataset using different feature types and numbers of clusters.

	w	w-e	w-e-h	t	t-h	h	e	e-h	e-t	e-t-h
500	215.05	273.26	236.81	8.06	22.61	29.68	276.83	113.0	93.86	73.06
1000	206.01	254.71	219.79	8.03	22.43	29.28	255.5	103.9	88.49	67.94
3000	200.38	238.57	205.31	8.01	22.26	28.88	240.08	95.46	84.04	63.38
5000	200.04	234.42	202.34	8.0	22.19	28.81	237.57	93.67	83.76	62.53

Table 5. Quality of estimated models (BBC 30, NYT 125 and Abstracts 1000) in terms of model precision and topic log odds.

Dataset	Model Precision	Dataset	Topic Log Odds
Total	0.87	Total	-4.23
BBC	0.93	BBC	-2.77
NYT	0.87	NYT	-4.57
Abstracts	0.8	Abstracts	-5.35

Table 6. Quality of estimated models in terms of model precision and topic log odds for every feature combination.

Feature Set	Model Precision	Feature Set	Topic Log Odds
w-e-h	1.00	e-t-h	-2.52
w-e	0.97	e	-2.58
e	0.97	t	-3.14
e-t-h	0.90	t-h	-3.90
h	0.87	h	-4.41
t-h	0.83	w	-4.76
t	0.83	w-e	-4.81
w	0.80	e-t	-5.04
e-h	0.77	e-h	-5.11
e-t	0.73	w-e-h	-6.05

Next, we calculated and evaluated the model precision separately for every feature combination. The values can be found in Table 6 left. We were very happy to notice that the standard, bag-of-words based model, has been outperformed by seven enriched models, five of which being pure resource topic models. Also the fact that both remaining models based on words occupy first two places shows the meaningfulness of enriching topic models.

As a next step we took a look at topic log odds. Similarly to the case of entity intrusion, also in topic intrusion the baseline (bag-of-words based model) is outperformed. In this case only by five models out of nine but all of them are pure resource topic models. The values can be found in Table 6 right. These are very good outcomes that confirm our initial assumptions that incorporating background knowledge into feature vectors brings additional context and makes the topic model easier to understand.

Figure 1 depicts histograms of both used measures. Alongside the accumulated values for mp and tlo, we presented example topics and documents accounting for certain values. We can see that the topic about technological companies is easily separable from the entity London Stock Exchange, while University, as an intruder in a topic about fish species is harder to spot and “infiltrates” the topic better. On the right side, in the topic log odds histogram, we see that the document about rugby and the protest for a bad referee is easy to decompose into the topical mixture, while a DBpedia Abstract document about a 2007 Americas Cup sailing regatta is a very specific piece of text.

Apart from getting measurable, quantitative results from the intrusion-based evaluation we also told the survey participants we would be happy to receive feedback in a free form if they had some remarks about the tasks they saw. As expected in such cases, only a few (three) decided to make an extra effort and use this option. However, all three stated harmonically that often right answers to

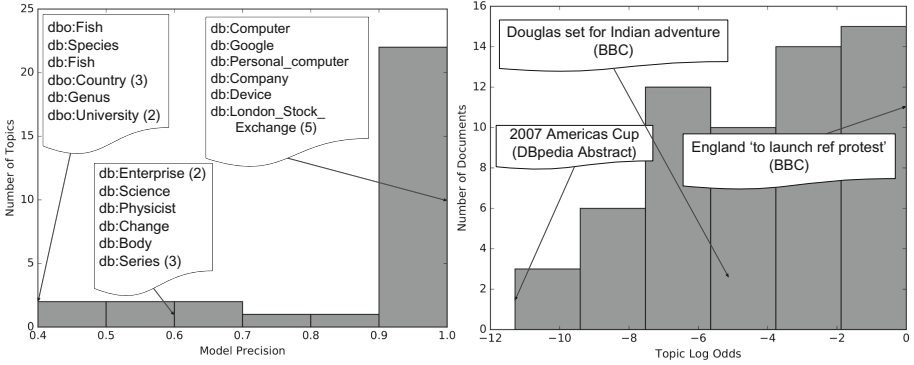


Fig. 1. Histograms of the model precisions of all 30 topics (left) and topic log odds of all 60 documents (right). On both histograms we presented several examples of topics and documents for the corresponding values, mp and tlo, respectively.

the topic intrusion tasks were guessed not by the descriptiveness of the presented text itself but rather by searching for an outlier within the group of four topics with no regard to the text.

To illustrate this phenomenon on an example, let's first take the following topic intrusion task, generated from the model mined on NYT Annotated Corpus dataset:

Of the 35 members in the Millrose Hall of Fame, all but one are renowned athletes. The only exception is Fred Schmertz, the meet director for 41 years.

Now, Schmertz will be joined by his son, Howard, 81, who succeeded him and directed the Millrose Games, the world's most celebrated indoor track and field meet, for 29 years, until 2003.

The 100th Millrose Games, to be held Friday night at Madison Square Garden, will be the 73rd for Howard Schmertz, now the meet director emeritus. He will be honored tomorrow night as this year's only Hall of Fame inductee.

FRANK LITSKY

SPORTS BRIEFING: TRACK AND FIELD

The generated answers can be seen below. The intruder is marked.

- dbo:City, dbo:Settlement, dbo:AdministrativeRegion, dbo:Person, dbo:Office Holder, dbo:Company, dbo:Town, dbo:EthnicGroup
- dbo:City, dbo:Settlement, dbo:AdministrativeRegion, dbo:Person, dbo:University, dbo:Disease, dbo:Town, dbo:Magazine
- dbo:Person, dbo:Film, dbo:Settlement, dbo:City, dbo:Magazine, dbo:Company, dbo:TelevisionShow, dbo:Newspaper
- **dbo:Country, dbo:PoliticalParty, dbo:MilitaryUnit, dbo:Settlement, dbo:Person, dbo:OfficeHolder, dbo:Weapon, dbo:MilitaryConflict**

Now let us explain the thought process that decided about selecting the right answer. These are our observations confirmed by the survey participants. First,

we see that two first topics are very similar which makes it very unlikely that one of them is an intruder. Several of third topic's top words are present either in the top words of the first (dbo:City, dbo:Company) or the second (dbo:Settlement, dbo:City, dbo:Magazine) topic, plus the ontology class Newspaper seems related to class Magazine in human understanding. Following this reasoning we conclude that the last topic is the intruder as it contains entities related to politics/military strategy. Unfortunately, this strategy is the only way to choose a correct answer for this task because all four answers seem unrelated to the presented document.

While evaluating and trying to understand the results, we came across another topic intrusion task from the NYC dataset which is a perfect example of how counter-intuitive and uninterpretable the topics can be:

BASEBALL

American League TEXAS RANGERS–Acquired RHP Brandon McCarthy and OF David Pisanso from the Chicago White Sox for RHP John Danks, RHP Nick Masset and RHP Jacob Rasner.

National League MILWAUKEE BREWERS–Agreed to terms with RHP Jeff Suppan on a four-year contract.

FOOTBALL

National Football League MINNESOTA VIKINGS–Agreed to terms with DT Kevin Williams on a seven-year contract extension. PITTSBURGH STEELERS–Signed LB Richard Seigler from the practice squad. Released WR Walter Young.

HOCKEY

National Hockey League (...)

The possible answers can be seen below. Number of votes per answer can be seen in brackets. The intruder is marked.

- db:Agent, db:Preparation, db:Subtype, db:Epidemic, db:Birds, db:State, db:Agency, db:Fowl (2)
- db:System, db:Company, db:Software, db:Computer, db:Product, db:Application, db:Device, db:Transmission (3)
- db:Team, db:Player, db:Goaltender, db:Disk, db:Tournament, db:Position, db:Trophy, db:Hockey (0)
- **db:Club, db:Footballer, db:Sport, db:Cup, db:League, db:Team, db:Competition, db:Player** (0)

To our giant surprise, we noticed that the intruder is the last answer – one of two topics that we would connect with sports at all. This example clearly shows what a difficult task it is to interpret associations between documents and enriched topics.

5 Conclusion and Future Work

The results we achieved are very satisfying and promising for the area of topic models enriched with background knowledge which is still relatively new and

has not been subject to much open research so far. In our experimental setup we showed that our approach outperforms the established bag-of-words based models, i.e. the enrichment step we perform and injecting linked data while generating the feature vectors makes the topic models more understandable for humans. This is a very meaningful result since interpreting and labeling topics is an interesting and popular research area. Having topics containing resources linked to knowledge graphs would allow for completely new possibilities in this domain.

Automatic topic labeling would not be the only research and application area where we think enriched topic models could be used. We imagine scenarios where we apply enriched topic modeling, attach huge text corpora to knowledge bases and thus have systems to automatically “classify” unseen documents and to find their places as mixtures of subsets of the knowledge graph.

Assessing the quality of hierarchical topic models was out of scope for this work. However, if we define a measure for goodness of topical hierarchies and manage to estimate enriched models that satisfy certain condition (e.g. perform similarly well as their bag-of-words counterparts), we could attach mined topical hierarchies to knowledge bases. This way we could conduct a fully automatic taxonomy mining from arbitrary document collections and thus offer support for experts in every domain.

Naturally, the evaluation didn’t bring only positive outcomes. We encountered some issues that hinder unlocking the full potential of our approach. First of all, a vital obstacle while considering enriched models and pure resource models is the length of documents. This problem plays an even bigger role in our approach than in pure word-based models. When texts are too short and thus the NER systems fail to deliver a reasonable amount of resources, the inferencer is unable to predict the coverage proper in terms of human understanding because the feature vector is far too sparse. To overcome this issue, we could try to mine a model using Wikipedia dump, or at least a representative and reasonable subset of it, to have a universal, multi-domain topic model. Running a NER system, such as DBpedia Spotlight, on Wikipedia articles (and maybe experimenting with tuning its parameters, such as support and confidence) could result in much denser feature vectors of higher quality.

Second, we find that the fact that topic models mined using DBpedia Abstracts perform worse might, apart from the average document length, be influenced by the number of topics in the model. In this work we focused on the aspects of feature engineering and feature selection, i.e. defining new features from background knowledge and investigating which subset of them would perform best for a fixed number of topics. More patience when selecting k ’s and tailoring them better for every dataset should bring improvement in the quality of topic models.

Another important thing to note is that we didn’t touch the topic modeling algorithm itself. We only changed ways a document is represented as feature vector. Digging deeper into LDA in order to differentiate between words and

resources and maybe applying graph-based reasoning would be another idea worth experimenting with in the future in pursuit of *good* enriched topic models.

One more aspect we consider worth mentioning and discussing is the philosophical nature of this work. We were using human judgment to assess the quality of topic models. Chang et al. concluded that human judgment does not correlate to established information-theoretic measure of topic models. However, human judgment itself is never uniquely defined and depends on the knowledge, cultural background, etc. It is very subjective and humans can easily discuss and argue over a topic, its interpretation and usefulness.

Unfortunately, performing very thorough evaluation of this kind was beyond our possibilities. Generated tasks, especially topic intrusion tasks, were very time consuming. Not only they required thorough readings of the presented text but also precise elaborations of the answers. As we already mentioned, completing the full questionnaire (60 tasks of each type) took almost four hours. Our vision and a future direction to overcome this issue would be to arrange a setup with much more topics (several k 's per dataset) and, given appropriate funding, use Amazon Mechanical Turk⁶, a platform where human workers are paid a small cash reward for each completed task. Not only would it allow us to measure topic interpretability on a bigger, more representative group of subjects but also, as already mentioned, to examine more models. Performing such an extended evaluation could confirm, which dataset produces better topics – as for now we only know that BBC performed best in this given setup but maybe k 's we selected were not optimal for the chosen datasets. Also, we used a particular subset of the New York Times Annotated Corpus for the evaluation. Maybe our choice was unfortunate and lowered the quality of estimated models.

Results achieved in this work justify enriching topic models with background knowledge. Even though rather basic, our approach of extending the bag-of-words showed potential that enriched topic models impose. Digging deeper into the algorithm and differentiating between words and resources could further improve the quality of estimated models. That being said, given more research, interest and funds, in near future the enriched topic models could develop into more sophisticated methods applied for numerous purposes in multiple domains.

References

1. Blei, D.M., Ng, A.Y., Jordan, M.I.: Latent dirichlet allocation. *J. Mach. Learn. Res.* **3**, 993–1022 (2003)
2. Chang, J., Gerrish, S., Wang, C., Boyd-Graber, J.L., Blei, D.M.: Reading tea leaves: how humans interpret topic models. In: *Advances in Neural Information Processing Systems*, pp. 288–296 (2009)
3. Gabrilovich, E., Markovitch, S.: Feature generation for text categorization using world knowledge. *IJCAI* **5**, 1048–1053 (2005)
4. Garla, V.N., Brandt, C.: Ontology-guided feature engineering for clinical text classification. *J. Biomed. Inf.* **45**(5), 992–998 (2012)

⁶ <https://www.mturk.com/>.

5. Hoffman, M., Blei, D.M., Bach, F.: Online learning for latent Dirichlet allocation. In: *Advances in Neural Information Processing Systems*, vol. 23, pp. 856–864 (2010)
6. Hu, Z., Luo, G., Sachan, M., Xing, E., Nie, Z.: Grounding topic models with knowledge bases (2016)
7. Mimno, D., Wallach, H.M., Talley, E., Leenders, M., McCallum, A.: Optimizing semantic coherence in topic models. In: *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pp. 262–272. Association for Computational Linguistics (2011)
8. Newman, D., Chemudugunta, C., Smyth, P.: Statistical entity-topic models. In: *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 680–686. ACM (2006)
9. Newman, D., Lau, J.H., Grieser, K., Baldwin, T.: Automatic evaluation of topic coherence. In: *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pp. 100–108. Association for Computational Linguistics (2010)
10. Pinoli, P., Chicco, D., Masseroli, M.: Latent Dirichlet allocation based on Gibbs sampling for gene function prediction. In: *2014 IEEE Conference on Computational Intelligence in Bioinformatics and Computational Biology*, pp. 1–8. IEEE (2014)
11. Scott, S., Matwin, S.: Feature engineering for text classification. *ICML* **99**, 379–388 (1999)
12. Todor, A., Lukasiewicz, W., Athan, T., Paschke, A.: Enriching topic models with DBpedia. In: Debruyne, C., et al. (eds.) *OTM 2016. LNCS*, vol. 10033, pp. 735–751. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-48472-3_46
13. Wallach, H.M., Murray, I., Salakhutdinov, R., Mimno, D.: Evaluation methods for topic models. In: *Proceedings of the 26th Annual International Conference on Machine Learning*, pp. 1105–1112. ACM (2009)
14. Zong, W., Feng, W., Chu, L.-K., Sculli, D.: A discriminative and semantic feature selection method for text categorization. *Int. J. Prod. Econ.* **165**, 215–222 (2015)



Predictive Quality: Towards a New Understanding of Quality Assurance Using Machine Learning Tools

Oliver Nalbach^(✉), Christian Linn, Maximilian Derouet, and Dirk Werth

AWS-Institute for Digitized Products and Processes gGmbH,
Saarbrücken, Germany

{oliver.nalbach,christian.linn,maximilian.derouet,
dirk.werth}@aws-institut.de

Abstract. Product failures are dreaded by manufacturers for the associated costs and resulting damage to their public image. While most defects can be traced back to decisions early in the design process they are often not discovered until much later during quality checks or, at worst, by the customer. We propose a machine learning-based system that automatically feeds back insights about failure rates from the quality assurance and return processes into the design process, without the need for any manual data analysis. As we show in a case study, this system helps to assure product quality in a preventive way.

Keywords: Data analytics · Machine learning · Neural networks
Quality assurance · Preventive quality · Predictive quality

1 Introduction

Product defects are dreaded by manufacturers. Not only do they come with significant costs for replacement products, shipping and customer care staff, a high rate of defects will also damage a manufacturer's reputation. The longer it takes to detect a defect in the development and production process, the more expensive undoing the resulting damage becomes (Fig. 1, orange line). At the same time, quality management research [2] shows that most product failures can be traced back to early decisions in the product design process where far-reaching decisions are made (Fig. 1, yellow line). Unfortunately, negative consequences do not show up until the product is actually produced (e.g., in quality checks) or worse, until it is returned by the customer (Fig. 1, green line). To minimize expenses, insights about failure rates have to be fed back from the quality assurance and return processes into the design process. However, such a preventive failure avoidance is not common [2]. Manufacturers try to minimize defect rates in the production process but a causal connection to design flaws is not drawn, partly because those are rarely obvious.

The original contribution of this paper is a reconsideration of what the ideal properties of modern quality assurance should be and how those can be achieved

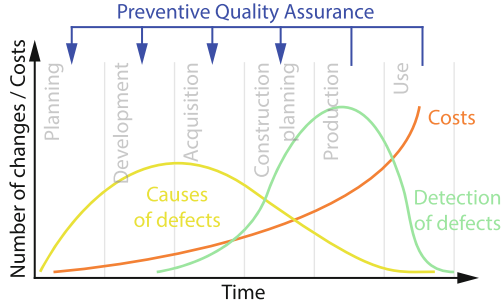


Fig. 1. Defect causes and detection, according to DGQ Beratung GmbH [1]. (Color figure online)

using machine learning tools. We do not only propose a concept but a complete system architecture for *Preventive Quality Assurance* (PreQA). It comprises two components, one for automated data analysis and another for real-time assistance of the product development process. The analysis component (Sect. 4.2) leverages data collected in quality assurance and the return process to identify correlations between product features and defect probabilities automatically using supervised learning. The assistance component (Sect. 4.3) can augment software such as design tools or ERP systems and inform their users in case an increased risk for product flaws is detected, feeding back analysis results (Fig. 1, blue arrows). Additionally, it offers help in resolving problems by suggesting alternatives to achieve a balance between freedom of choice and product quality.

After reviewing previous work (Sect. 2), we introduce our concept of Preventive Quality Assurance, distinguishing it from the current understanding of preventive and predictive quality (Sect. 3), and go on to propose an architecture for its realization (Sect. 4). By a case study in clothing industry (Sect. 5) we demonstrate the potential of PreQA even in creative industries with their vast freedom of decision. We conclude with an outlook towards future work on the topic (Sect. 6).

2 Previous Work

The current use of machine learning in manufacturing focuses on the fabrication itself [3]. Historical data is analyzed to optimize future fabrication processes and to reduce defect rates. According to recent studies [4] these analytical methods bear potential for vast cost reductions but consider only part of the production pipeline. PreQA instead can combine data from multiple steps of the pipeline, i.e., quality checks at different stages of the process as well as customer feedback to automatically gain insights which then can be used in the earliest design stages.

Preventive quality management in enterprises is governed by manual processes. One example is Failure Mode and Effects Analysis (FMEA) [5] in

which failure modes are identified and rated regarding their severity, probability of occurrence and probability of detection. The product of these ratings is called Risk Priority Number (RPN) and used to order the failure modes in terms of urgency. Another example is Fault Tree Analysis (FTA) [6] where the conditions for a particular fault to occur are mapped to variables in a logic formula. The value of the expression then indicates whether a particular configuration of conditions may cause the fault or not. Both FMEA and FTA are based on experience and subjective judgment of the people involved which makes them time consuming and unreliable. We instead perform a fully autonomous analysis which is unbiased and does not require the intervention of human experts.

Statistical process control (SPC) [7] is a semi-automatic way to realize quality control using statistics. After manual identification of quality-relevant quantities, those are monitored, usually in an automated fashion, using tools such as control charts or histograms. This way, the variation of a manufacturing processes can be tracked and its violation of defined bounds can be noticed easily. Going further, the reactions following steps undertaken to decrease variability become immediately visible. As a whole, the tools of SPC can therefore aid an iterative quality refinement process but still require a lot of human interaction, unlike the preventive quality assurance system proposed in this paper.

Due to its manual nature, preventive quality management is not yet common. In particular not in creative industries such as the clothing industry. According to a recent study [2], 74% of textile manufacturers do not use any kind of quality management methods when developing new products. The intent of PreQA is to change this by automation and simplification, offering a way to achieve quality assurance even in trades with a diverse and creative development process.

So far, little approaches exist to support product design by software-based assistance systems. Schulte and Pape [8] propose an assistance tool embedded into the CAD system used by a designer supporting her in complying with a formal definition of the design process. While such an approach might increase overall efficiency, it does not offer recommendations how to prevent product failures based on actual data. Dienst et al. [9] present a prototypical system which analyzes information collected throughout the life cycle of a product in order to improve future generations of the same product, e.g., by replacing machine parts known to be error-prone by higher quality parts of the same type. As such, the system can only be used for iterative improvement of a specific product, while PreQA offers predictions for entirely new product designs which may differ drastically from any existing product for which data has been collected.

In some industrial branches, an iterative quality improvement can be achieved using control systems. At the core of a (closed-loop) control system, a controller tries to adjust the input to a pre-defined system such that its output approaches an ideal value. The challenge lies in simultaneously accounting for the effects that the system has on the controller itself. In case the controlled system is a manufacturing device and the output value is directly related to quality of the produced item, e.g., in wafer production [10], this yields a production environment which autonomously and rapidly stabilizes the quality of its output at a

high level. The scope of such approaches is limited to selected areas and cannot easily be transferred, for example to clothing industry (Sect. 5) where the input to the manufacturing process is high-dimensional and the delay between design decisions (control actions) and feedback (returns) can be a year or even longer.

3 A New Understanding of Quality Assurance

The ISO9000 standard for quality management [11] defines *quality assurance* as the “part of quality management focused on providing confidence that quality requirements will be fulfilled”. The future tense in this definition indicates that quality assurance has to be proactive: adverse effects of quality flaws have to be prevented before they become relevant in the actual use of a product, e.g., by rigid definition of how a product is to be assembled. A closely related term is predictive quality which refers to methods that, based on data, identify statistical patterns to foresee future developments concerning the quality of a product. For example, even if an automotive part like a steel spring has passed quality controls meeting the primary expectations, the actual wear and tear in practice may be outside the range initially assumed during product design entailing negative consequences. The only way to circumvent such problems used to be the extremely conservative specifications leading to partly unnecessary expenses. Nowadays, cars contain a multitude of different sensors using which data about the real-world use of a vehicle can be captured. A manufacturer can compare the real use to the initial assumptions and react before urgency arises thanks to predictive methods.

In our opinion, quality assurance should be both predictive, that is, draw conclusions from data, and preventive, i.e., resolve causes of problems, not symptoms. Quality assurance should not just predict that problems with a car’s suspension will arise to schedule timely recalls. And quality assurance should not simply prevent that flawed products slip through to the customers, still generating waste. Ideally, the flaws themselves should be prevented by eradicating their causes which are often found in unfavorable design decisions. The tool to achieve this is prediction based on data but the prediction has to happen before a product is even produced, not after. Furthermore, this process should not have to rely on expensive, time consuming and error prone manual analysis but should be enabled by automated machine learning. In the following section, we will outline an architecture achieving just that.

4 Reference Architecture for Preventive Quality Assurance

We will now describe the general architecture and functionality of a Preventive Quality Assurance system, also termed PreQA in the following, from a theoretical, formal point of view and independently from a particular application area. Results from an actual implementation of the system within the scope of a case study in clothing industry will follow in Sect. 5.

4.1 Overview of the System

PreQA consists of an autonomous analysis and a real-time assistance component. The assistance component provides real-time feedback to product designers or managers regarding their decisions while the analysis component generates the data structures necessary to provide the assistance functions.

The input to PreQA consists of heterogeneous sources of information about products and their (implicitly associated) failure rates, e.g., data from enterprise resource planning (ERP) software, technical documents detailing product designs or quality check logs. These structured data may be enriched by unstructured data like photographs of the defective products or textual descriptions. In general, the prerequisite to PreQA is information about existing products, their features (cf. Sect. 4.2) and cases of defective products.

The output of the system, provided by the assistance component, consists of two functions: First, the system judges the current state of a product and notifies the person responsible if an increased risk for defects is detected, showing possible defects with their likelihood and samples of similar existing products exhibiting the defects to serve as a justification and to increase acceptance of the system. Second, the system can be queried for alternatives which are as similar as possible to a defined configuration but have a smaller estimated failure rate.

4.2 Automated Data Analysis

The data analysis component learns a mapping from products to associated defect rates based on information collected for existing products. In particular, this mapping can then be applied to entirely new and unseen products, too.

Preliminary Definitions. We define *products* as vectors whose elements correspond to product features f_i from the set $F = \{f_i \mid i \in \{1, \dots, n_{\text{feat}}\}\}$ of all features. Features can be nominal, i.e., discrete and unordered, or real-valued: $F = F_{\text{nom}} \dot{\cup} F_{\text{real}}$. Nominal features f have associated (non-empty and finite) sets of values R_f while $R_f = \mathbb{R}$ for $f \in F_{\text{real}}$. As we also want to run predictions for partially-defined products such as they appear at different stages of the design and manufacturing process, we define products over feature subsets: For a selection I of features, let $S_I = \times_{i \in I} R_{f_i}$ the set of (possible) products exhibiting feature set I . Then, the complete set of products is

$$P = \left\{ p \mid p \in S \in \left\{ \times_{i \in I} R_{f_i} \mid I \subseteq \{1, \dots, n_{\text{feat}}\} \right\} \right\}.$$

As an example consider $n_{\text{feat}} = 3$ with $R_{f_1} = \{\text{shirt}, \text{jacket}\}$, $R_{f_2} = \{\text{India}, \text{China}\}$, $R_{f_3} = \mathbb{R}$. Feature f_1 specifies a type of garment, f_2 the country of origin and f_3 may be a yarn weight. Then $\{\text{shirt}, \text{China}, 52.3\} \in R_{f_1} \times R_{f_2} \times R_{f_3}$, $\{\text{India}\} \in R_{f_2}$ and $\{\text{shirt}, 61.4\} \in R_{f_1} \times R_{f_3}$ are valid products.

Crucial to PreQA is the prediction of defect probabilities: Let D the set of defects, $d \in D$ a particular defect, e.g., a ripped seam at a shirt's collar,




Fig. 2. A sample generic hierarchy of possible defects that might arise in a product.

and I a feature selection. Then by $c_{I,d}^* : S_I \rightarrow [0, 1]$ we denote a hypothetical function which, given values for the features in I , outputs the probability that d occurs. PreQA approximates $c_{I,d}^*$ by fitting a *classifier* $c_{I,d}$ to sample data of the form $(\bar{p}, \bar{d}) \in P \times \{d\} \dot{\cup} \{\text{none}\}$, i.e., pairs of products and a second value which either declares the defect $d \in D$ that occurred or that there was none. We call a collection $m_I = \{c_{I,d} \mid d \in D\}$ a *meta classifier* for feature set I with the $c_{I,d}$ being *sub-classifiers*. As information about defects d may be more or less specific, e.g., the location of a defect may be known or not, we define a partial order on D using a tree (Fig. 2) grouping defects hierarchically. We call d' *at least as specific* as d , denoted $d' \leq d$, if d' is in the subtree that has d as root.

Analysis Flow. The upper part of Fig. 3 shows the analysis flow. In a preliminary step, structured input is pre-processed (Fig. 3a) and transformed to a representation (Fig. 3b) which allows sampling of vectors (\bar{p}, \bar{d}) as defined above. Unstructured data associated with individual samples of structured data is mined for details on the type and location of the respective defect (Fig. 3c). If the identified defect d' is more specific than the known d from the structured data, the latter is overridden (Fig. 3d).

The training of new meta classifiers may be initiated by two events. When initializing the system, we may train a base set of meta classifiers, e.g., for all feature selections $\{i\}$ with $i \in \{1, \dots, n_{\text{feat}}\}$ containing just one feature (Fig. 3e). Second, the assistance system may schedule new meta classifiers, too, if a good classifier for a particular input product is missing (Fig. 3f). When a new meta classifier m_I is requested we train the corresponding sub-classifiers (Fig. 3g). To train $c_{I,d}$, we need relevant training and test data and, for this, select samples (\bar{p}, \bar{d}) with $\bar{p} \in S_J$, $I \subseteq J$ and $\bar{d} \in \{d' \leq d\} \dot{\cup} \{\text{none}\}$, i.e., samples where the product information covers all features in I and the defect is at least as specific. After m_I has been trained, we rate its quality (Fig. 3h), e.g., by the average of the area under the receiver operating characteristic curve (ROC AUC) metrics [12] for all sub-classifiers, and store m_I together with its rating $r(m_I)$ (Fig. 3i).

To provide a justifying evidence for warnings issued by the assistance (Sect. 4.1), we use a k-nearest neighbor (knn) query in the set of defective samples selecting the k most similar existing products for which a defect previously occurred. Fast knn searches require appropriate data structures like k-d trees [13]. As each search is specific to one sub-classifier—only neighbors regarding

the specific feature subset and with a relevant defect are considered—we trigger the construction of these data structures after training a new meta classifier (Fig. 3j).

4.3 Real-Time Assistance

The real-time assistance can answer two kinds of queries, one regarding risks involved in a particular product design and a second one to retrieve alternatives with higher predicted quality. In this section, we explain how these queries can be dealt with conceptually. How and when they are triggered, e.g., automatically in the background or manually by the user, depends on the software which is augmented by PreQA and is not part of this paper.

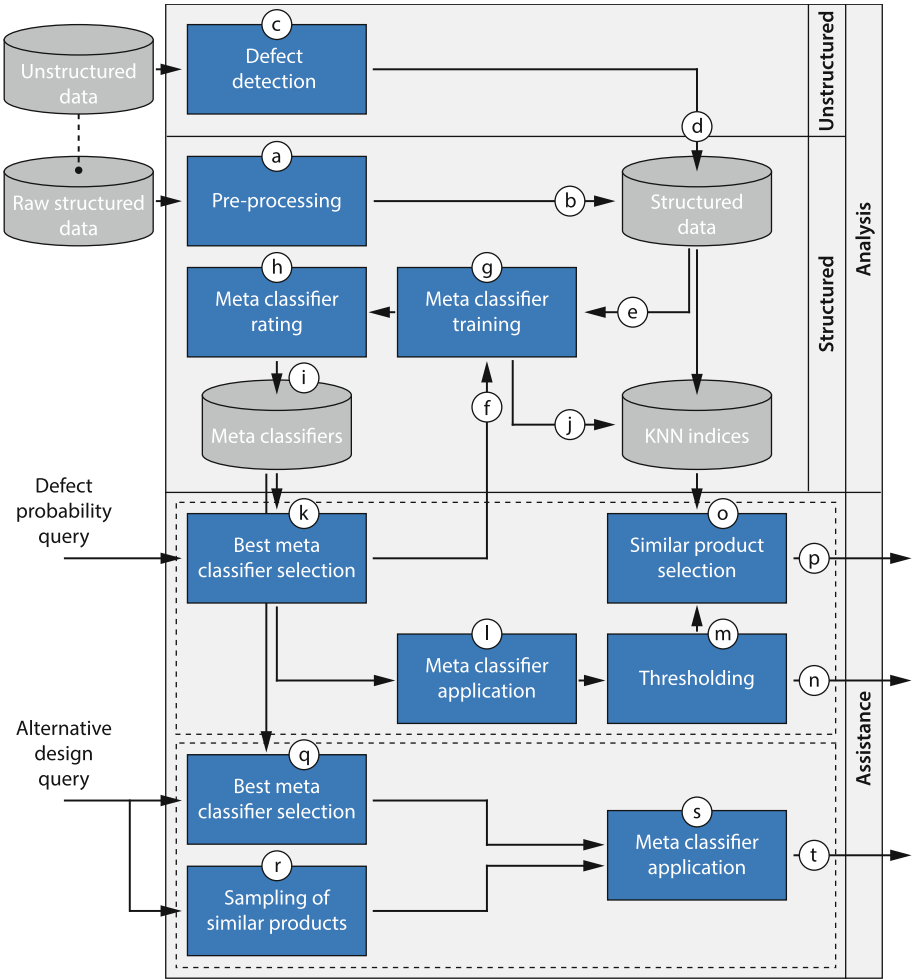


Fig. 3. Flow diagram of the analysis and assistance components.

Defect Probability Queries. A query for potential issues consists of a product p with features I . We retrieve all applicable meta classifiers m_J with $J \subseteq I$ and select the one with the highest rating

$$m_I^* = \operatorname{argmax}_{m_J, \text{ s.t. } J \subseteq I} r(m_J),$$

which we expect to be the most reliable one (Fig. 3k). When m_I^* uses only a subset of features ($J \neq I$), we schedule training of m_I in the background (Fig. 3f) to enable better predictions in the future. Then, we apply all sub-classifiers in m_I^* to p to estimate the probability of individual defects (Fig. 3l) and perform a thresholding (Fig. 3m) which filters out unlikely defects to not bury users of the system in warnings. If no defect passes the thresholding the response to the query is empty. Otherwise pairs of defects and their estimated probabilities are returned (Fig. 3n) together with the k most similar products determined by the knn searches associated with the relevant sub-classifiers (Fig. 3o+p).

Suggestion of Alternative Designs. The input for a query of less defect-prone alternatives consists of a product p with features I and a user-defined subset $N \subseteq I$ of attributes whose values should stay the same for suggested alternatives. We select the best meta classifier m_I^* with feature subset $J \subseteq I$ as before (Fig. 3q). It may happen that the user has locked all attributes used by m_I^* , i.e., $N \supseteq J$. In this case, we return an empty response to state that no alternative has been found. Otherwise, we sample alternative products p' (Fig. 3r) meeting the user's conditions ($p'_i = p_i$ if feature $p_i \in R_f$ for $f \in N$), apply m_I^* to all of them (Fig. 3s), and return the k samples with the smallest cumulative defect prediction (Fig. 3t). How intelligently samples can be picked depends on the nature of the features. A general approach is to only modify one feature value of p at a time, sampling all possible values for nominal features and picking a constant number of equidistant values for scalar features.

5 Case Study: PreQA in Clothing Industry

We are currently evaluating PreQA in a case study on data from a German clothes manufacturer. For this, we have implemented a prototype (Fig. 6) which currently supports the analysis of structured input data. Integration of unstructured data and implementation of the real-time assistance are planned for the near future.

5.1 Input Data

Input data has been extracted from an ERP system used by the clothing manufacturer. This includes information about individual products and their features, e.g., the type of garment, material mix or country of origin, but also production volume data and a set of return transactions containing the individual ID of the affected product, case-specific information, for example regarding

sizing, and the defect that occurred. The specialized defect catalog is defined manually as the ERP system does not define a hierarchy itself. In total, we deal with a dataset containing 120k product variants and information about 150k return transactions.

5.2 Implementation Details

Software. Our implementation of PreQA is based on Python for its platform independence and wide range of data science libraries. In particular we use scikit-learn [14] for machine learning algorithms and pandas [15] for data processing.

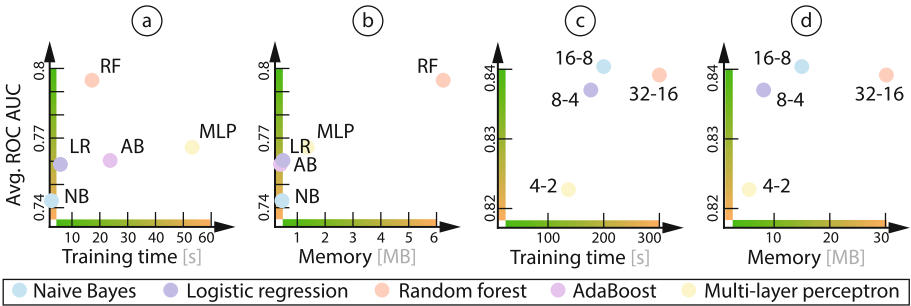


Fig. 4. Comparison of classification methods regarding time (a) and memory consumption (b) and of MLP layouts (c+d). Different methods (a+b) were compared on a typical feature subset, while MLP layouts were compared on the full feature set to estimate an upper bound on the necessary complexity.

Classification Method. We evaluated different classification algorithms for implementing meta classifiers: logistic regression, naïve Bayes, random forests, multi-layer perceptron-style neural networks (MLPs) and AdaBoost based on decision trees. MLPs offer the overall best trade-off between precision, speed and memory consumption. More precisely, we use MLPs with two hidden layers containing 16 and 8 hidden units, respectively, and sigmoid activation functions. This architecture was chosen after comparing a set of possible network layouts (Fig. 4c+d).

Data Pre-processing. The extracted product information (Sect. 5.1) is pre-processed in three steps before we can use it to train MLPs. In a first step, constant attributes, which usually correspond to unused columns of the ERP database, are removed automatically (Fig. 5a \rightarrow b). In the second step, some nominal attributes are translated to (vectors of) scalars to incorporate prior knowledge. As an example, the ERP system mined by us encodes material mixes as strings (e.g., “80% cotton, 20% polyester”) which would restrict prediction

to existing mixes. Instead, we parse the strings and add real-valued features in the range $[0, 1]$ encoding the contribution of each material (Fig. 5b \rightarrow c). In the last step, we label-encode remaining nominal features (Fig. 5c \rightarrow d) using the LabelEncoder class provided by sci-kit learn [14].

Sampling of Training Data. A peculiarity of our setting are its severely imbalanced classes. Flawless items outweigh defective ones by several orders of magnitude. We down-sample the non-defective class, typically by a factor of $\gamma = 10^{-2}$. Down-sampling the negative class leads to over-estimation of defective cases. While we could accommodate for this by scaling the output of the trained MLP [16], we are only interested in whether the prediction of the MLP is above a certain threshold (Fig. 3m). Thus, we instead simply adjust our thresholds accordingly. Before the samples are fed into the network, we apply one-hot encoding to the nominal features to avoid unwanted order relations.

Training Details. We train our MLP classifiers using the Adam algorithm [17], applying a log-loss. Training is automatically terminated once the loss does not change by at least $5 \cdot 10^{-4}$ for two consecutive iterations. The sub-classifiers are trained in parallel worker threads. As an example, training a typical meta classifier including 15 sub-classifiers on 170-element inputs (after one-hot encoding) requires about 50 s in four parallel threads on an Intel Xeon Gold 5122 server CPU.

(a)

id	material	designer	country	weight
0	20%C, 80%P	NULL	IND	0.23
1	50%P, 50%PA	NULL	CN	0.1
2	100%C	NULL	CN	0.31
3	100%W	NULL	ITA	0.07
4	30%W, 70%P	NULL	GER	0.53

(b)

id	material	country	weight
0	20%C, 80%P	IND	0.23
1	50%P, 50%PA	CN	0.1
2	100%C	CN	0.31
3	100%W	ITA	0.07
4	30%W, 70%P	GER	0.53

(c)

id	C	P	PA	W	country	weight
0	0.2	0.8	0.0	0.0	IND	0.23
1	0.0	0.5	0.5	0.0	CN	0.1
2	1.0	0.0	0.0	0.0	CN	0.31
3	0.0	0.0	0.0	1.0	ITA	0.07
4	0.0	0.7	0.0	0.3	GER	0.53

(d)

id	C	P	PA	W	country	weight
0	0.2	0.8	0.0	0.0	0	0.23
1	1.0	0.5	0.5	0.0	1	0.1
2	1.0	0.0	0.0	0.0	1	0.31
3	0.0	0.0	0.0	1.0	2	0.07
4	0.0	0.7	0.0	0.3	3	0.53

Fig. 5. Pre-processing consists of feature removal, feature translation and label encoding.

Interface to Other Software. We have chosen to design a general system which is supposed to connect with arbitrary other software. To make this possible, we have included a server component which handles requests in JavaScript Object Notation (JSON) [18] for the use cases described in Sect. 4.3. For example, products can be sent to the server as JSON objects mapping feature names to strings or float values, defect probabilities are returned as objects mapping defects (specified as strings) to floating point probability values.

5.3 Preliminary Results

While the implementation of assistance functions and their integration with third-party software is still in progress, analysis results can already be queried by manual application of meta classifiers. In Fig. 6 we show the output of a meta classifier processing products specified by the type of garment, their cotton and polyester percentages, size and country of origin. As can be seen, PreQA identifies relations ranging from trivial facts, such as that shirts do not have zippers (6b), to expert knowledge, e.g., that polyester knitwear is prone to pilling (6a), to observations which otherwise could only be made using expensive manual data analysis, for example the trade-offs between different countries of production (6b). Regarding computational performance, training a meta classifier typically requires not much longer than a minute (Fig. 4a) and thousands of meta classifiers can be stored simultaneously in the memory of a conventional server for fast selection (Fig. 4b). The queries yielding the results shown in Fig. 6 take less than 10 ms to process.

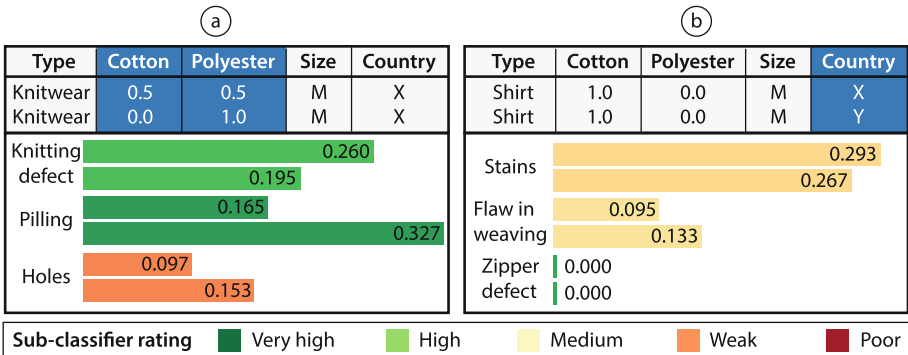


Fig. 6. Analysis results comparing design alternatives. Left: two knitted garments of different materials. A higher percentage of polyester reduces the risk of knitting defects at the cost of increased pilling. Right: the same shirt produced in different countries shows varying risks of stains and weaving flaws. As expected, there is no risk of defect zippers as they are not common in shirts. Note that probabilities are overestimated because we typically down-sample the significantly larger class of flawless items by a factor of about 100 due to performance reasons.

6 Conclusion and Outlook

We have presented the idea of Preventive Quality Assurance, a system guiding people responsible for developing new products in their decisions in order to optimize product quality. Its unique feature is the degree of automation: a user does not need to be an expert in data analytics but is automatically provided with the knowledge and assistance she needs at every point in time. PreQA has not only been outlined conceptually (Sect. 3) but also formally (Sect. 4) and has been demonstrated as a promising approach in a case study (Sect. 5).

Future work on PreQA will, first, complete the implementation of the system. For this, aspects such as the analysis of unstructured data, the realization of assistance functions and an integration with other software will have to be dealt with. Combining multiple sources of data beyond ERP systems alone, such as product design files or quality checks would help to increase the precision of predictions and the applicability in more steps of the development process.

Acknowledgments. This work is based on *Preventive Quality Assurance*, a project partly funded by the German ministry of education and research (BMBF), reference number 01S17012D. The authors are responsible for the publication's content.

References

1. FMEA - Fehlermöglichkeits- und Einflussanalyse. Whitepaper, Deutsche Gesellschaft für Qualität (2012)
2. Morlock, S.: Qualitätsgesicherte Produktentwicklung in der Bekleidungsindustrie (2013). https://www.hohenstein.de/de/inline/pressrelease_45248.xhtml?excludeId=45248
3. Harding, J.A., Shahbaz, M., Srinivas, S., Kusiak, A.: Data mining in manufacturing: a review. *Manuf. Sci. Eng.* **128**, 969–976 (2006)
4. Veronesi, L., Kifonti, R., Cattaneo, G.M.: Data-driven innovation in the manufacturing industry. *Eur. Data Market Study* (2014)
5. Stamatis, D.H.: Failure Mode and Effect Analysis: FMEA from Theory to Execution. ASQ Quality Press, Milwaukee (2003)
6. Ericson, C.A.: Fault tree analysis. In: Hazard Analysis Techniques for System Safety, pp. 183–221 (2005)
7. Oakland, J.S.: Statistical Process Control. Routledge, New York (2007)
8. Schulte, R., Pape, D.: Advanced digital mock-up based on the usage of assistances. *Des. Principles Pract.* **5**, 195–206 (2011)
9. Dienst, S., Fathi, M., Abramovici, M., Lindner, A.: Development of a knowledge-based feedback assistance system of product use information for product improvement. *Int. J. Prod. Dev.* **19**, 191–210 (2014)
10. van der Meulen, S.H., Tousain, R.L., Bosgra, O.H.: Fixed structure feedforward controller design exploiting iterative trials: application to a wafer stage and a desk-top printer. *Dyn. Syst. Meas. Control* **130**(5), 051006 (2008)
11. Quality management systems - fundamentals and vocabulary. Standard, International Organization for Standardization (2015)
12. Bradley, A.P.: The use of the area under the ROC curve in the evaluation of machine learning algorithms. *Pattern Recogn.* **30**(7), 1145–1159 (1997)

13. Bentley, J.L.: Multidimensional binary search trees used for associative searching. *Commun. ACM* **18**(9), 509–517 (1975)
14. scikit-learn. <http://scikit-learn.org>
15. Python data analysis library (pandas). <https://pandas.pydata.org>
16. Dal Pozzolo, A., Caelen, O., Johnson, R.A., Bontempi, G.: Calibrating probability with undersampling for unbalanced classification. In: 2015 IEEE Symposium Series on Computational Intelligence, pp. 159–166 (2015)
17. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization. arXiv preprint [arXiv:1412.6980](https://arxiv.org/abs/1412.6980) (2014)
18. Javascript object notation (JSON). <https://www.json.org>

Business and Enterprise Modelling



Application of Inductive Reference Modeling Approaches to Enterprise Architecture Models

Felix Timm^(✉), Katharina Klohs, and Kurt Sandkuhl

Chair of Business Information Systems, University of Rostock,
Rostock, Germany

{felix.timm, katharina.klohs,
kurt.sandkuhl}@uni-rostock.de

Abstract. Enterprise architectures (EA) help organizations to analyze interrelations among their strategy, business processes, application landscape and information structures. Such ambitious endeavors can be supported by using reference models for EA. In recent years, research has increasingly been investigating inductive methods for reference model development. However, the characteristics of EA models have not been considered in this context yet. We therefore aim to adapt existing inductive approaches to the domain of reference enterprise architectures development. Using design science research our work contributes to the reference modeling community with (i) a comparative analysis of inductive reference modeling methods regarding their applicability to EA models, (ii) the refinement of an identified approach to reference EA design and (iii) its application in a use case.

Keywords: Reference enterprise architecture · Inductive reference modeling
Enterprise architecture · Literature review

1 Introduction

Enterprises need to be aware of the relations among their strategy, processes, applications and infrastructures to be able to rapidly react on changing demands in the market and within their organization. The Enterprise Architecture (EA) research domain contributes to this purpose by providing methods and tools to establish a more holistic perspective on enterprises [1, 2]. This includes to systematically capture and develop an EA using modeling languages like ArchiMate [3]. Such models represent different architectural layers of an enterprise, such as business, application and technology architecture. Since EA projects are highly time- and resource-consuming, organizations would benefit from reference models for EAM. A Reference Enterprise Architecture (R-EA) can be defined as a generic EA for a class of enterprises that is used as foundation in the design and realization of the concrete EA [4].

Many methods exist to develop reference models—both deductive and inductive approaches [5]. Especially, the latter one gains increasing attention by the research community in recent years [6]. Unfortunately, inductive approaches focus primarily on business process model structures [7] and research lacks investigating their applicability towards EA structure like ArchiMate [8]. Nevertheless, we claim that the benefits

inductive methods provide for reference process development will also qualify for R-EA development once these methods can be adapted.

In order to fill this research gap, we propose a method for inductive derivation of a R-EA using approaches of existing inductive methods. We therefore use design science research (Sect. 2) in the domain of inductive reference modeling (Sect. 3). We perform a literature review on current inductive methods and assess them regarding their applicability to EA modeling against prior defined requirements (Sect. 4). We further refine one suitable approach to the EA domain (Sect. 5) and demonstrate the new method by dint of a use case (Sect. 6).

2 Research Design

We structure our research report in terms of the design science research (DSR) methodology proposed by [9]. Therein we aim to extend existing methods for developing inductive reference process model to their application to the domain of R-EA development. Peffers et al. define five activities for DSR projects. We performed them as follows [9]:

For the *problem identification*, we conducted a literature study to investigate the existing body of knowledge in the field of inductive reference modeling methods and their existence in the EA domain (see Sects. 3 and 4). We combined the approaches by Kitchenham [10] and Webster and Watson [11] to thoroughly study existing research. The results verified our stated absence of an inductive approach that addresses EA structures. In the DSR activity to *define the objectives for a solution* we analyzed EA structures and defined requirements they possess towards inductive approaches (see Sect. 4.1). Using these, we performed a comparative analysis and identified suitable approaches and chose the most appropriate one to adapt it for the development of a R-EA, which was the activity to *design and development* in our DSR process (see Sect. 5). The fourth activity is *demonstration* that the artifact actually solves the problem. We therefore applied the resulting artefact to eleven individual EA models and developed a R-EA that fulfilled the prior defined requirements (see Sect. 6). The last DSR activity, *evaluation*, has the intention to observe and measure how well the artifact offers the solution to the given problem. We discuss the benefits and drawbacks of the refined artefact and derive future research activities from these insights (see Sect. 7).

3 Inductive Reference Modeling

In general, reference models are information models developed for a certain problem in a certain application domain. The purpose of their development is to be reused in a concrete application case in this domain. The reuse of a reference model is intended to increase both efficiency and effectivity of an enterprise's information systems and their change management [12]. From a user-oriented perspective, Thomas understands a reference model as a model used to support the construction of another model [13]. From the perspective of reusability, other authors such as vom Brocke argue that reference models are characterized by the concepts of universality and recommendation

[12]. The life cycle of such reference models can be distinguished between the phase of construction and the phase of application [5, 14]. By presenting insights in reference enterprise architecture development we contribute to the first phase, i.e. the construction of reference models.

For constructing reference models, research discusses two generic strategies. While the deductive reference modeling derives reference models from generally accepted knowledge, the inductive approach abstracts from individual models to agree on a common understanding within the reference model [15]. Most established reference models have been developed based on deductive approaches [16]. However, inductive reference modeling provides potential because more and more relevant data in terms of logs and concrete information models of organizations are available. Further, inductively developed reference models tend to have a higher degree of detail, are more mature and seem to be more accepted when it comes to reference model application [6].

Fettke provides a seven phase method for inductive reference model construction [17]. The core of the approach is the derivation of a reference model from presorted individual information models. Although Fettke names clustering as a technique to elicit the reference model, it is not made explicit how it is conducted. In recent years, reference modeling related research addresses this topic and approaches were developed. One work applies abstraction techniques from business process mining to the reference model domain [7]. Other approaches utilize for example natural language processing techniques [18], graph theory [19] or clustering methods [20], but always focus on business process model structures.

4 Comparative Analysis of Inductive Reference Modeling Approaches

In this section we systematically investigate inductive methods regarding their applicability to EA model structures. Therefore, Sect. 4.1 defines criteria, which were used to examine, whether an inductive approach is considered applicable to EA model structures. Section 4.2 presents the results of our systematic literature review and thus the identified approaches before we identify suitable approaches in Sect. 4.3.

4.1 Requirements for Inductive R-EA Development

Through EA models the complex interrelations between an enterprise's organizational and operational structure with used information systems, processed data and realizing technologies are made explicit. Such models consist of layers and elements, which define different perspectives on the enterprise [2]. We use the TOGAF framework [21] as it is widely accepted among practitioners and comes with a modeling language ArchiMate in version 3.0 [3].

In order to derive requirements for the selection of suitable approaches we shortly distinct business process and EA models from each other. Ahlemann et al. state that EA models are more comprehensive than pure business process models since they represent an organization from different perspectives and are not restricted to the business layer [1]. Greefhorst and Proper define five elements of an EA [22]: *concerns*, which

are related to EA stakeholders and group distinct interests on the EA model; architecture *principles* that guide the EA model; an EA *model*, which relates different elements of the EA with each other; *views* that represent a projection of the EA model regarding a certain concern; and, a *framework* that provides a meta model and modeling guidelines. Further, EA models result in more complex model structures. For example, event-driven process chains (EPCs) define six elements and three relation types while ArchiMate's core framework defines 36 elements and 12 relation types. In summary, business process models can be interpreted as a partial model of EA models, since they can be integrated into an EA as Lankhorst et al. demand [2].

REQ1: The Approach Should Be Suitable to Other Modeling Languages. The authors of the approach have to mention the applicability to other modeling languages. This might be other process modeling languages or any modeling language. We dismiss the approach if the authors state that it is solely applicable to the utilized model structure. Otherwise, we concluded this by studying the approach in detail.

REQ2: The Approach Does Not Solely Rely on the Control Flow of the Model Structure. Most inductive approaches focus on the control flow characteristic of business process languages like ECP to derive a reference model. Next to such dynamic relationships, EA models use many other relationship types like the assignment (structural), serving (dependency) or specialization relationship. Thus, we dismiss approaches that are not applicable to static model structures.

REQ3: The Approach Has to Be Adjustable Regarding Concepts of Viewpoint and Concerns of ArchiMate. EA models focus on a more aggregated level of detail in contrast to business process models, which represent a much higher level of detail [2]. ArchiMate uses different viewpoints to address different concerns of the model's stakeholders [22]. The approach may use mechanisms to vary the decision what elements to integrate into the reference model depending on a viewpoint's concern. For example, more criteria than a frequency threshold could be used. This also might support the identification of best practices among the individual models.

We consider any approach that fulfills all three requirements suitable for the application of R-EA development. However, we do not directly dismiss approaches fulfilling REQ1 and REQ2 but not REQ3, because it may still be applicable to sub-models of the ArchiMate models.

4.2 Results of the Literature Review

For a sound systematic literature review (SLR), research questions should be defined that underlie the review process [10]. The aim of our SLR was to identify articles that propose a method or an approach how to derive inductively a reference model from a set of individual models. These have to address the same problem domain in the same modeling language and, thus, can be somehow processed to derive a common practice or a best practice among them. Further, we wanted to analyze these approaches, whether they can be applied to EA models by using the REQs presented in Sect. 4.1.

After several alterations we used the search term (reference NEAR/4 model*) AND (inductive* OR mining). It was performed on the document

title, abstracts and keywords and we used a NEAR-operator in order to gather result like “reference process models”. Furthermore, we included the term “mining” since related work often use it instead of “inductive” [23]. The query was performed on the databases SCOPUS (296 results), AISeL (20 results) and IEEEExplore (94 results) from 2007 until today. This period was chosen since prior work in [6] already conducted a similar search, which we wanted to extend with the most recent approaches. The authors only identified inductive approaches after 2007 in their work. Then, we excluded irrelevant articles by reading all titles and abstracts. This resulted in 16 results, from which we eliminated three duplicates and then reviewed the 13 full documents. We included every article in English and German language that explicitly addressed inductive reference modeling and propose an approach for reference model derivation. We excluded work that offered techniques, e.g. from the process mining domain, but lacked in describing its application to reference model development. Additionally, we used related literature reviews to include relevant articles that we did not capture yet [6, 7, 23, 24]. In the end, we identified 21 approaches for inductive reference modeling. We analyzed them against REQ1, REQ2 and REQ3. The results are presented in Table 1.

Table 1. Identified approaches and their fulfilment of requirements. (“+” fulfils REQ; “-” does not fulfil REQ; “o” may be applicable with major adjustments)

APPROACH	REQ1	REQ2	REQ3	DECISION	APPROACH	REQ1	REQ2	REQ3	DECISION
Rehse et al. [25]	-	-	-	Inapplicable	Li et al. [26]	-	o	-	Inapplicable
Scholta [27]	+	+	+	Applicable	Ling, Zhang [18]	+	o	-	Dismissed
Leng, Jiang [28]	+	-	-	Dismissed	Li et al. [29]	-	o	-	Inapplicable
Ardalani et al. [16]	+	+	-	Applicable	La Rosa et al.[30]	+	o	-	Dismissed
Rehse, Fettke [31]	+	-	-	Dismissed	Fettke [20]	-	o	-	Inapplicable
Li, et al. [19]	-	o	-	Inapplicable	Yahya et al. [32]	+	o	-	Dismissed
Gottschalk et al.[33]	-	o	-	Inapplicable	Rehse et al. [7]	+	o	-	Dismissed
Rehse et al. [34]	-	o	-	Inapplicable	Aier et al. [35]	+	o	-	Dismissed
Martens et al. [36]	+	+	-	Applicable	Li et al. [37]	-	o	-	Inapplicable
Sonntag et al. [38]	-	+	-	Inapplicable	Yahya et al. [39]	+	o	-	Dismissed
Martens et al. [40]	+	o	-	Dismissed					

4.3 Identification of Suitable Approaches

Only three approaches were identified that are suitable for their application to EA models regarding our requirements, i.e. the approaches from [16, 27, 36]. Although only the approach by Scholta [27] fulfilled all three requirements we deem the approaches by Ardalani et al. [16] and Martens et al. [36] applicable to EA models as well. Nine approaches were assessed as inapplicable. Their authors explicitly state that they were specifically developed for certain model structures like EPCs [34], WS-BPEL [26] or Workflow Nets [33]. Some focus on the analysis of process models’ control flow [25, 37] and, thus, cannot be applied to EA model structures. The

remaining approaches were dismissed due to major adjustments one would have to make using them for EA model structures. For example, some use graph theory to derive reference models [32], while others use clustering algorithms [30] or transformation matrices for their approach [40]. Although it is possible to represent ArchiMate models in graphs or matrices, the approaches' effort highly increases with the complexity of different relationships ArchiMate defines. We exclude them from the scope of this work, but discuss their potential for future research at the end of this article.

Based on the defined requirements in Sect. 4.1 we chose the approach by Scholta [27] for applying it to ArchiMate models. Scholta's is the only approach that proposes a method for inductively deriving a best practice reference model and therefore defines a number of characteristics that also can be applied to ArchiMate models, which fulfils REQ 3. Still, in a prior research paper we already successfully applied the approach by Ardalani et al. [16] to ArchiMate models by manually calculating the minimal graph-edit distances of ArchiMate model parts. The same may be done for the approach by Martens et al. [36]. The remainder of this paper depicts our application of Scholta's approach and demonstrates the results by dint of a use case.

5 An Adjusted Approach for Inductive R-EA Development

The approach by Scholta aims to semi-automatically derive reference process models that represent best practices from a set of individual process models [27]. The approach is called "RefPA" and addresses the problem domain of public administration. However, we understand it to be applicable to other problem domains as well, although the author omits to discuss "RefPA's" generalizability. The fact, that the approach, to date, is only conceptually described allows us to modify it on the conceptual level, since the procedure is well documented.

RefPA defines the following objectives: to keep the source models' structure; to identify commonalities of the source models; to consistently group certain segments of the source models; and, to evaluate these groups in order to identify best practices. After collecting all source models in step 0, Scholta defines five steps to develop the reference model. In step 1 all source models are merged together, containing all nodes and edges of any source model. Step 2 subsequently detects all nodes that occur in all source models and, hence, are common elements. In step 3 Scholta uses the concept of the SQL constructs GROUP BY, WHERE and CONTAINS to identify groups in each source model. For instance, process nodes are grouped by a certain document they process (GROUP BY), while they may fulfil certain conditions (WHERE) or a concrete node may have to be in this group (CONTAINS). After groups were identified, they are evaluated and ranked using various criteria. Scholta lists 21 different criteria, such as processed documents, lead-time or personnel costs. Depending on the criteria and the problem domain, the best groups are identified as best practices. Finally, in step 5 all common elements are nodes as well as the best groups are assembled together to a reference models, which contains common as well as best practices, as Scholta claims. In the following, we will discuss these five steps regarding necessary adjustments for ArchiMate models. For more insights into the approach we encourage the reader to study the work by Scholta in [27].

Step 0: Provide Source Models. In contrast to process models, ArchiMate models may be of high complexity. Hence, we define further requirements towards the source models: (i) likewise to [27], we demand every source model to address the same problem domain; (ii) each source model has to comply with a predefined viewpoint structure, which assigns each viewpoint used to a certain EA concern according to [2]—otherwise the source model needs to be aligned to it; and, (iii) in order to guarantee consistent identification of commonalities and best practices, source models have to represent the same level of granularity.

Step 1: Create Merged Model. For each viewpoint from the model structure we conduct one iteration of Step 1 → Step 2 → Step 3 → Step 4. We understand ArchiMate viewpoints as sub models, since they represent different concerns regarding the EA model [2]. The following adjustments address the reference model derivation for one single viewpoint before they are assembled to the resulting reference model in Step 5. A Merged ArchiMate Model is created, using one model repository. Each source model is represented by the corresponding ArchiMate viewpoint. This guarantees a merged model, while the source model's varieties are still documented by single viewpoints. This helps conducting step 2. As a preparation, we create a reference viewpoint for the viewpoint at hand. If the source viewpoints follow a common structure of elements on the lowest level of granularity, it should be integrated in the reference model as well.

Step 2: Identify Common Elements. In line with Scholta, we analyze the source viewpoints for common elements and relations. To date, we assume that elements with the similar semantic are labeled with the similar syntax. In addition, we integrate a frequency threshold as used in [16]. It may vary for each viewpoint and is defined depending on the reference model's purpose. In contrast to Scholta, Ardalani et al. define an element a common practice if it occurs in the majority of source models [16]. A prior defined threshold is applied for each element and relation of the source viewpoint. All elements and relations above the threshold are integrated in the reference viewpoint.

Step 3: Group Elements. Although this is essential to derive best practices, it is the vaguest step of Scholta's approach. Again, ArchiMate models are much more complex than process models and, thus, there are a lot different possibilities to group model parts. Still this does not guarantee the identification of best practices. During the case study (see Sect. 6) we identified approaches for grouping the viewpoints' elements. This depends from the EA layer, which the current viewpoint addresses. On the business layer, sub-functions or sub-process may be grouped as well as processes that are assigned to the same business role. On the application layer, data objects may be grouped regarding the phenomenon they relate to (e.g. customer profile). Again, the overall purpose of the reference model has to be considered before grouping elements. For each source viewpoints the same grouping decision has to be applied.

Step 4: Evaluate Groups. We deem all criteria proposed by Scholta to be suitable to ArchiMate models. Next to them, there may exist domain-specific criteria, depending on the problem the reference model intends to address. Consequently, before choosing the criteria one has to understand how the model at hand may help the future model

users to solve this problem in the best way. Moreover, it has to be clarified whether the available source models offer the required information. The ranking of element groups may use multiple criteria. For different viewpoints, different criteria may be applied.

Step 5: Assemble Reference Model. After all reference models are developed, the reference viewpoints are integrated into the final reference model. There may exist duplicates, since some elements or relations are defined in various reference viewpoints.

6 Case Study from the Financial Domain

We applied the adjusted RefPA approach to R-EA development in the financial domain. It aims to holistically capture all relevant aspects of the financial organization affected by regulation in order built a reference model that helps financial institutes to effectively and efficiently implement a compliance organization. Therefore, the reference model follows the structure of enterprise architectures and uses. One part of the reference model addresses the prevention of “other criminal acts” (abbr. ssH for German “sonstige strafbare Handlungen”), which is regulated in §25 h Abs. 1 KWG and applies for every financial institution [41]. Such crimes may be fraud, corruption or treason.

For this part of the reference model, we conducted and transcribed 11 structured interviews with responsible persons from distinct financial institutes. We developed the structure of the interviews using deductive techniques. Therefore, we consulted compliance experts for a first structure of a ssH prevention system and studied the KWG law as well as guidelines provided by the Federal Financial Supervisory Authority [42]. Afterwards, all interviews were transferred to EA models using the same modeling structure and guidelines, which was done according to Lankhorst [2]. Each individual EA model was structured by eight different ArchiMate viewpoints, which displayed different aspects of the EA models. Each viewpoint was related to a certain purpose. In the following, we concentrate on one specific viewpoint in order to demonstrate the application of Scholta’s RefPA approach. As depicted in Sect. 5 the procedure was applied for each viewpoint.

In *Step 0* all requirements were fulfilled: we clarified the model’s purpose (i); we defined a thorough viewpoint structure (ii); and, since we used the same structured interviews the same level of granularity of the source models could be assumed (iii). We conducted the subsequent steps for each of the eight ArchiMate viewpoints. This is demonstrated by dint of one viewpoint that addresses how prevention of other criminal acts is conducted on a strategic level. The viewpoint’s purpose is to recommend financial institutes what business functions should be implemented on a management level for preventing such crimes and further, what business roles are related to these tasks. For this purpose, we use the Business Function Viewpoint, which is defined by The Open Group [3]. In *Step 1* a merged model was created, which incorporated all business functions and roles as well as their relations among each other. We defined a corresponding reference viewpoint next to the source viewpoints. As a structural framework for the reference viewpoint we used the general management framework we

deduced during the interview design—each interviewee verified it. These functions were namely “Building an Organizational Framework”, “Corporate Hazard Analysis”, “Implementation of a Prevention System” and “Monitoring”.

In *Step 2*, we identified common elements and relations in the viewpoint. After a CSV-Export we defined a threshold of 80% and identified 41 common business functions, and five common business roles. We argue that using this threshold still enables to capture commonalities, while we can ignore outliers that may occur due to the interview as a model elicitation method. We further assumed that elements with the same semantic were labeled with the same syntax because we used the same model element library when modeling the source models. During conducting *Step 3*, we decided to group the source viewpoint’s models by dint of the four generic business functions identified during Step 1 (see Fig. 1). Group G3 splits into two groups, namely “G3.1 Measures for Internal Hazards” and “G3.2 Measures for External Hazards”. The intention was to gather not only the common business function related to the four management functions, but also to accompany them with identified best practices from the source viewpoints. To do this evaluation in *Step 4* we found it most reasonable to use a criterion that identifies the group, which contains most aspects mentioned by the Federal Financial Supervisory Authority in [42]. In the case that multiple groups from different source viewpoints met this criterion, we integrated the group that contained the most business functions. During Step 5 we integrated the groups from the prior steps and assembled the reference viewpoint.

Figure 1 depicts an extract of the developed reference viewpoint for implementing a prevention system for other criminal acts on a strategic level. The complete view has 96 business function elements, which are assigned to eleven distinct business roles via

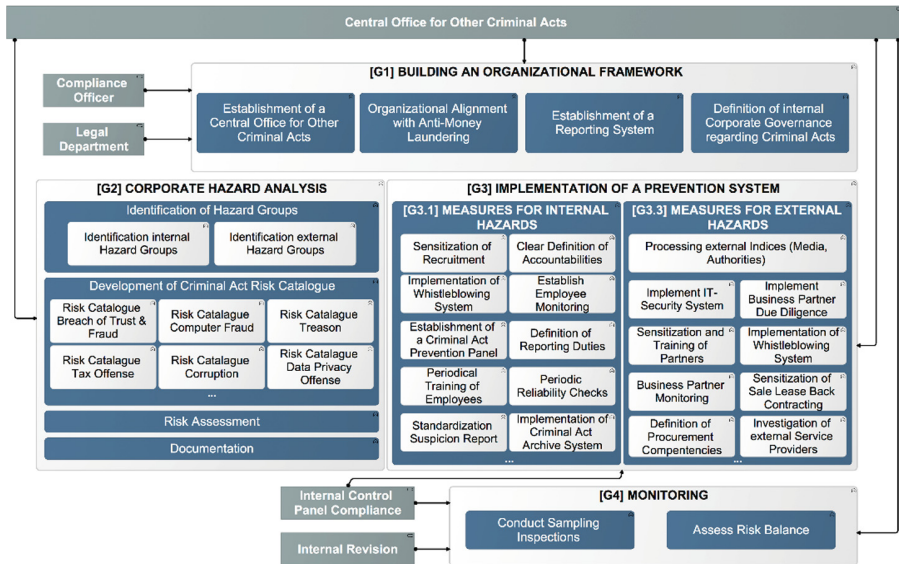


Fig. 1. Reference viewpoint: strategic level of preventing other criminal acts

16 assignment relations. Further, 95 composition relations are included. The Open Group defines a *Business Function* as an internal behavior and groups the behavior based on resources, skills, knowledge, etc. It is performed by a *Business Role* (using the *assignment* relation), which is defined as the responsibility to perform specific behavior and requires certain competencies from individuals to qualify for it. Composition among business functions is represented by dint of the *composition* relation [3]. In Fig. 1 visualizes this by graphical composition among the business functions.

7 Discussion of Results and Conclusion

In this article we address the research gap that the reference modeling discipline lacks an inductive approach tailored for EA models. By applying a DSR process we therefore define requirements that EA models have towards an inductive derivation of reference enterprise architecture (R-EA) based on multiple individual EA models. We apply these to existing inductive methods that we identified during a systematic literature review, and which almost merely address business process structures. We identify the approach by Sholta [27] to be the most appropriate for our purposes and refine it in order to demonstrate its successful application in a use case for R-EA development. Domain experts validated the resulting R-EA in a first workshop session. Nevertheless, there are several open issues observed for both (i) the developed artefact in particular and (ii) inductive approaches for R-EA development in general we identify for future research.

First, the current approach is conducted manually using ArchiMate modeling tools and spreadsheets. Once the individual EA models meet the input requirements of Step 0, the following steps may be automatized—including the grouping of Step 3. Second, the best practice elicitation of Step 4 needs to be investigated in more detail. Although the experts in the validation phase considered our chosen criteria as plausible, we see flaws in it. Since the individual models were gathered before, may have been certain information we could have inquired, if we would have known what criteria we would use in Step 4 of the approach. Further, the concept of best practice is very vague and should be investigated in much more detail for appropriate criteria selection. Last, when automatizing the approach, it may be interesting to move from a viewpoint-based approach to a model library based approach. This would allow deriving the reference model at once and preventing the effortful model integration.

From a general perspective towards the inductive development of a R-EA, some of the ideas and algorithms of the identified approaches from the literature review should still be considered for further adaptation. First, EA models could be transformed into graph structures in order to automate the merging step; as done in [30, 39]. Second, clustering methods could be applied in order to identify commonalities beyond single EA elements but also sub-models; as done in [18, 39]. Third and last, natural language process could help to analyze syntactical and semantical similarities of EA elements, that are not labeled similarly, but describe the same phenomenon [30, 32].

References

1. Ahlemann, F., Stettiner, E., Messerschmidt, M.: Strategic enterprise architecture management. In: Ahlemann, F., Stettiner, E., Messerschmidt, M., Legner, C. (eds.) *Challenges, Best Practices, and Future Developments*. Springer, Heidelberg (2012). <https://doi.org/10.1007/978-3-642-24223-6>
2. Lankhorst, M.: Enterprise architecture at work. In: *Modelling, Communication and Analysis*. Springer, Heidelberg (2017). <https://doi.org/10.1007/978-3-642-29651-2>
3. ArchiMate 3.0 Specification: Open Group Standard (2016)
4. ten Harmsen van der Beek, W., Trienekens, J., Grefen, P.: The application of enterprise reference architecture in the financial industry. In: Aier, S., Ekstedt, M., Matthes, F., Proper, E., Sanz, Jorge L. (eds.) *PRET/TEAR -2012*. LNBIP, vol. 131, pp. 93–110. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-34163-2_6
5. Fettke, P., Loos, P.: Referenzmodellierungsforschung. *Wirtschaftsinf* **46**, 331–340 (2004)
6. Rehse, J.-R., Hake, P., Fettke, P., Loos, P.: Inductive reference model development: recent results and current challenges. In: Mayr, H.C., Pinzger, M. (eds.) *INFORMATIK 2016*, 26–30 September 2016, Klagenfurt, Austria, vol. P-259. GI, Bonn (2016)
7. Rehse, J.-R., Fettke, P., Loos, P.: Eine Untersuchung der Potentiale automatisierter Abstraktionsansätze für Geschäftsprozessmodelle im Hinblick auf die induktive Entwicklung von Referenzprozessmodellen. In: Alt, R., Franczyk, B. (eds.) *WI2013*, Leipzig, Germany (2013)
8. Timm, F., Sandkuhl, K., Fellmann, M.: Towards a method for developing reference enterprise architecture. In: Leimeister, J.M., Brenner, W. (eds.) *WI2017*, St. Gallen (2017)
9. Peffers, K., Tuunanen, T., Rothenberger, M.A., Chatterjee, S.: A design science research methodology for information systems research. *J. Manag. Inf. Syst.* **24**, 45–77 (2007)
10. Kitchenham, B., Pretorius, R., Budgen, D., Pearl Brereton, O., Turner, M., Niazi, M., Linkman, S.: Systematic literature reviews in software engineering – a tertiary study. *Inf. Softw. Technol.* **52**, 792–805 (2010)
11. Webster, J., Watson, R.T.: Analyzing the past to prepare for the future: writing a literature review. *MIS Q.* **26**, 13–23 (2002)
12. Vom Brocke, J.: *Referenzmodellierung. Gestaltung und Verteilung von Konstruktionsprozessen*. Logos, Berlin (2003)
13. Thomas, O.: Understanding the term reference model in information systems research: history, literature analysis and explanation. In: Bussler, C.J., Haller, A. (eds.) *BPM 2005*. LNCS, vol. 3812, pp. 484–496. Springer, Heidelberg (2006). https://doi.org/10.1007/11678564_45
14. Thomas, O.: *Management von Referenzmodellen. Entwurf und Realisierung eines Informationssystems zur Entwicklung und Anwendung von Referenzmodellen*. Logos-Verl., Berlin (2006)
15. Becker, J., Schütte, R.: Referenz-Informationsmodelle für den Handel: Begriff, Nutzen und Empfehlungen für die Gestaltung und unternehmensspezifische Adaption von Referenzmodellen. In: Krallmann, H. (ed.) *Wirtschaftsinformatik 1997*, Physica-Verlag HD (1997)
16. Ardalani, P., Houy, C., Fettke, P., Loos, P.: Towards a minimal cost of change approach for inductive reference model development. In: *ECIS* (2013)
17. Fettke, P.: Eine Methode zur induktiven Entwicklung von Referenzmodellen. In: *Tagungsband Multikonferenz Wirtschaftsinformatik 2014, MKWI 2014*, pp. 1034–1047 (2014)

18. Ling, J., Zhang, L.: Generating hierarchical reference process model using fragments clustering. In: Asia-Pacific Software Engineering Conference, vol. 2016, May 2016
19. Li, C., Reichert, M., Wombacher, A.: Discovering reference models by mining process variants using a heuristic approach. In: Dayal, U., Eder, J., Koehler, J., Reijers, H.A. (eds.) BPM 2009. LNCS, vol. 5701, pp. 344–362. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-03848-8_23
20. Fettke, P.: Integration von Prozessmodellen im Großen: Konzept, Methode und experimentelle Anwendungen. In: Thomas, O., Teuteberg, F. (eds.) Proceedings der 12. Internationalen Tagung Wirtschaftsinformatik (WI 2015), pp. 453–467. Osnabrück (2015)
21. TOGAF Version 9. Van Haren Publishing, Zaltbommel (2010)
22. Greefhorst, D., Proper, E.: Architecture Principles: The Cornerstones of Enterprise Architecture. Springer, Heidelberg (2011). <https://doi.org/10.1007/978-3-642-20279-7>
23. Rehse, J.R., Fettke, P.: Towards situational reference model mining - main idea, procedure model & case study. In: Leimeister, J.M., Brenner, W. (eds.) Proceedings der 13. Internationalen Tagung Wirtschaftsinformatik (WI2017), pp. 271–285. St. Gallen (2017)
24. Schoknecht, A., Thaler, T., Fettke, P., Oberweis, A., Laue, R.: Similarity of business process models—a state-of-the-art analysis. *ACM Comput. Surv.* **50**, 1–33 (2017)
25. Rehse, J.R., Fettke, P., Loos, P.: An execution-semantic approach to inductive reference model development. In: 24th ECIS (2016)
26. Li, C., Reichert, M., Wombacher, A.: Discovering reference process models by mining process variants. In: Proceedings of the IEEE International Conference on Web Services (2008)
27. Scholta, H.: Semi-automatic inductive derivation of reference process models that represent best practices in public administrations. In: 24th ECIS 2016 (2016)
28. Leng, J., Jiang, P.: Granular computing-based development of service process reference models in social manufacturing contexts. *Concurrent Eng.* **25**, 95–107 (2016)
29. Li, C., Reichert, M., Wombacher, A.: Mining business process variants. challenges, scenarios, algorithms. *Data Knowl. Eng.* **70**, 409–434 (2011)
30. La Rosa, M., Dumas, M., Uba, R., Dijkman, R.: Business process model merging: an approach to business process consolidation. *ACM Trans. Softw. Eng. Methodol.* **22**, 9 (2013)
31. Rehse, J.-R., Fettke, P.: Mining reference process models from large instance data. In: Dumas, M., Fantinato, M. (eds.) BPM 2016. LNBIP, vol. 281, pp. 11–22. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-58457-7_1
32. Yahya, B.N., Bae, H., Bae, J.: Generating valid reference process model using genetic algorithm. *Int. J. Innov. Comput. Inf. Control* **8**(2), 1463–1477 (2012)
33. Gottschalk, F., van der Aalst, Wil M.P., Jansen-Vullers, M.H.: Mining reference process models and their configurations. In: Meersman, R., Tari, Z., Herrero, P. (eds.) OTM 2008. LNCS, vol. 5333, pp. 263–272. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-88875-8_47
34. Rehse, J.R., Fettke, P., Loos, P.: A graph-theoretic method for the inductive development of reference process models. *Softw. Syst. Model.* **16**, 833–873 (2017)
35. Aier, S., Fichter, M., Fischer, C.: Referenzprozesse empirisch bestimmen. Von Vielfalt zu Standards. *Wirtschaftsinformatik & Management: WuM: die Praktikerzeitschrift für Wirtschaftsinformatiker* **3**, 14–22 (2011)
36. Martens, A., Fettke, P., Loos, P.: A genetic algorithm for the inductive derivation of reference models using minimal graph-edit distance applied to real-world business process data. *Tagungsband MKWI* **2014**, 1613–1626 (2014)

37. Li, C., Reichert, M., Wombacher, A.: The MinAdept clustering approach for discovering reference process models out of process variants. *Int. J. Coop. Info. Syst.* **19**, 159–203 (2010)
38. Sonntag, A., Fettke, P., Loos, P.: Inductive reference modelling based on simulated social collaboration. In: Leimeister, J.M., Brenner, W. (eds.) *WI2017*. St. Gallen (2017)
39. Yahya, B.N., Bae, H.: Generating reference business process model using heuristic approach based on activity proximity. In: *Smart Innovation, Systems and Technologies*, vol. 10 (2011)
40. Martens, A., Fettke, P., Loos, P.: Inductive development of reference process models based on factor analysis. In: Thomas, O., Teuteberg, F. (eds.) *Proceedings der 12. Internationalen Tagung Wirtschaftsinformatik (WI 2015)*, pp. 438–452. Osnabrück (2015)
41. Gesetz über das Kreditwesen. KWG (2017)
42. Bundesanstalt für Finanzdienstleistungsaufsicht (ed.): *Auslegungs- und Anwendungshinweise zu § 25c KWG (“sonstige strafbare Handlungen”)* (2011)



Towards a Typology of Approaches for Sustainability-Oriented Business Model Evaluation

Thorsten Schoormann^(✉), Anna Kaufhold, Dennis Behrens,
and Ralf Knackstedt

University of Hildesheim, Universitätsplatz 1, 31141 Hildesheim, Germany
{thorsten.schoormann, anna.kaufhold, dennis.behrens,
ralf.knackstedt}@uni-hildesheim.de

Abstract. While numerous studies related to business models are dealing with understanding and representing businesses, limited research focus on the evaluation aspect. However, due to the highly dynamic environment, changing availability of resources and booming digitalization, the evaluation of such models and alternatives is an essential task. Analysing, evaluating and innovating businesses is a multidimensional issue, which poses challenges that need to be facilitated. Thus, in this study, we aim to explore approaches and criteria for sustainability-oriented business model evaluation. Based on an extensive literature review, we obtained 50 articles and determined more than 40 different evaluation methods.

Keywords: Sustainability · Business model · Evaluation · Assessment

1 Introduction

The rapid deterioration of the natural environment as well as concerns over wealth disparity and corporate social responsibility present fundamental issues for our entire society [1]. Thus, ‘sustainability’ has increasingly gained importance in business research and practice [2, 3] and, although companies often pay attention to be economically competitive and effective, there is also an emerging stream that respects ecological and social issues, for example by trying to achieve cleaner production processes [4]. To foster this, approaches that support changes in behaviour and practice are required.

Here, methods and tools from business model research can contribute, which are usually applied to visualize, innovate and evaluate business models [5]. Although the evaluation of business models including their configuration options is emphasized by various authors (e.g., [4–6]), current research mainly addresses the representation [8–10], the categorization (e.g., via taxonomies) and the understanding or the relevance of such models [11–13]. However, evaluation is a fundamental aspect to facilitate fast decisions, which are essential today because of the increasingly dynamic and uncertain environment (e.g., through changing availability of resources, growing demands, incremental pollution or booming digitalization). Evaluation approaches should support

such fast and high quality decision-making and the adjustment of business strategies, for example to manage risks and uncertainty [14], analyse consequences of options [6, 15], or measure the performance of a business [16] also in respect sustainability.

Sustainability is usually divided into three dimensions, for example Triple Bottom Line (economic, ecological and social) [17]. According to the Brundtland report, it is the “development that meets the needs of the present without compromising the ability of future generations” [1]. Nonetheless, prior studies regarding business model evaluation tend to deal with economic-oriented assessment [18, 19] in term of aspects such as efficacy, robustness, profitability and feasibility [6, 11, 20–22]. Hence, we currently lack of approaches that respect ecological and social issues (e.g., emissions or impact on the society). Moreover, existing approaches are varying and are often focused on separated aspects because of different disciplinary backgrounds of the authors like management, accounting or IT [23]. Thus, we need an overview that structures existing approaches regarding multidimensional assessment strategies. This gap is problematic because it inhibits the management to make informed decision (e.g., selection of resources), manage their business model effects, improve their strategy in respect sustainable criteria and allocate future efforts in the right way. From an IS/tool designer perspective, it hinders to develop supporting software that facilitates these tasks. Accordingly, this paper is guided by the following question: *What are current approaches and criteria that allow for sustainability-oriented business model evaluation?*

In order to answer this, we carry out an extensive literature review and develop a typology of current methods, tools and criteria. In doing so, we aim to contribute to both practice (e.g., providing tools for managing sustainability-oriented effects) and research (e.g., design new tools and methods as well as derive advanced theories that, for example attempt to explain how specific methods affect the performance of business model evaluation). Following our research methodology (Sect. 3), we conduct a literature review, code the results and classify them in form of a typology. Next, we present our main findings, which contribute to the sustainability-oriented assessment, comparison, benchmarking of different business model aspects (Sect. 4). Finally, we discuss results, provide implications and conclude with our paper (Sect. 5).

2 Background

2.1 Business Models and Sustainability

In general, a business model “describes the rationale of how an organization creates, delivers, and captures value.” [8, p. 14] A good business model answers: Who is the customer? What does the customer value? How do we make money? How we can deliver value to customers at an appropriate cost? [24]. Regarding the term ‘business model’, some authors refer to (1) the way a company does business and others focus on (2) the model aspect while using the term ‘business model’ [25]. Our study follows the second stream in particular and understands the term in a conceptual manner.

The coherence of business models and sustainability has been increasingly discussed in research over the last few years (e.g., [7, 26–28]). Generally, a sustainable

business model seeks to capture “economic value while maintaining or regenerating natural, social, and economic capital beyond its organizational boundaries” [7, p. 1]. Because a shared and well-accepted understanding is still missing, most of the prior work focused on specific fields of application like mobility [29] or renewable resources [30]. Existing studies related to sustainability and business models mainly contribute to the representation or understanding of ecological and social issues (e.g., by extending current business model languages [9]). Thus, we currently lack comprehensive knowledge regarding a sustainability-oriented evaluation of such business models [19].

2.2 Evaluation of Business Models

The current literature mostly focuses on economic-oriented methods and criteria by, for example assessing quantitative aspects like profit, cash flows and costs [31]. Hamel [20] is one of the first authors who suggests evaluation approaches which basically consists of four key criteria, namely efficiency, uniqueness, fit and profit boosters. Afuah and Tucci [11] propose a quantitative-oriented method that allows for analysing profitability measures, profitability predictor measures and business model component attribute measures. In addition, Gordjin and Akkermans [13] develop an approach for evaluating the economic feasibility in a quantitative way (via profit sheets and what-if scenarios). Similar, Weill and Vitale [32] deal with performance measurements and critical success factors like cost savings per transaction. Besides focusing on financial issues, some authors provide more holistic methods, for example Osterwalder and Pigneur [8] suggest to apply a SWOT analysis and provide a set of key questions that could be used to improve businesses. Moreover, some studies deal with the evaluation of sustainability in particular. For instance, Lüdeke-Freund et al. [19] propose a conceptual framework for sustainability-oriented business model assessment by combining the Balanced Scorecard as a controlling tool and the Business Model Canvas.

Due to the fundamental importance of evaluation, some researchers already try to provide an overview of selected aspects. For example, Tesch und Brillinger [14] categorize methods for digital business model innovation, and Alexa [23] conceptualizes business model evaluation approaches from an entrepreneurial perspective. Regarding sustainability, Hansen et al. [33] propose a broad framework of approaches that focuses on assessing sustainability-oriented innovations, and Ness et al. [34] classifies environmental assessment methods, but do not focus on business models in particular. Nevertheless, existing studies often address only separated aspects of evaluation or provide overviews for specific domains (e.g., digital innovation), and a consolidated overview of approaches for sustainability-oriented business model evaluation is still missing.

3 Research Methodology

The primary goal of our study is to identify strategies for evaluating multidimensional sustainability in business models. Therefore, we conduct a literature review [35]. Because the researcher has to document literature findings, the selection of keywords and the evaluation of the results [36], we follow vom Brocke et al. [36] who allow for

methodological rigorousness. The authors define five steps: definition of review scope (Sect. 1), conceptualization of topic (Sect. 2), literature search (Sect. 4.1), literature analysis and synthesis (Sects. 4.2 and 4.3), and research agenda (Sect. 5).

4 Literature Search, Analysis and Synthesis

4.1 Literature Search

As a first step, we aimed to identify publications that apply, extend or develop an evaluation strategy for business models. Built upon an initial analysis of different search terms and their combinations, we specified the following search items: “*business model assessment*”, “*business model evaluation*” and “*sustainability*” or “*sustainable*”.

Within our search, we included various databases and search engines (*Google Scholar*, *Science Direct*, *Wiley Online Library* and *Taylor&Fancis Online*) to allow for an interdisciplinary and broad search. Moreover, we incorporated proceedings of high ranked *IS conferences (AISeL)* as well as high ranked *IS journals* by using the *AIS Seniors Scholars’Basket of Journals*. No further constraints (e.g., timespan) were applied.

Based on a search in 12/2017 and 01/2018, we found 521 articles (Fig. 1). As proposed by Webster and Watson [35], a complete keyword search, and an evaluation of titles and abstracts was applied to each article (Evaluation 1). In doing so, non-relevant articles were eliminated. The remaining articles (114) were verified by the full text (Evaluation 2). Finally, we removed duplicates (Consolidation). In total, we found 50 articles that propose evaluation approaches which addresses at least one dimension of sustainability. In order to contribute to the reliability and validity of the results as well the research process itself, two researchers evaluated the relevance of each article.

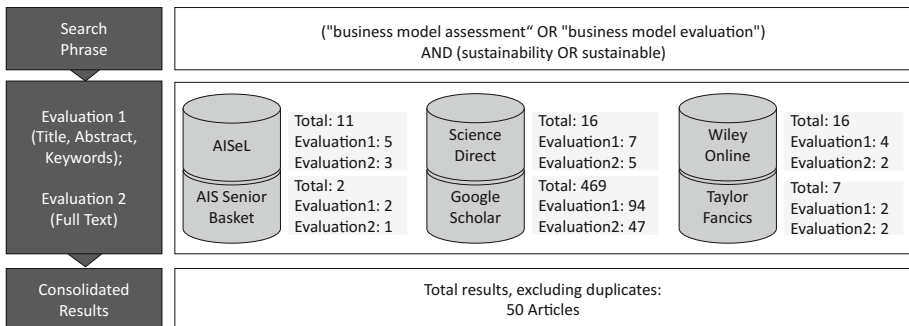


Fig. 1. Overall literature review process.

4.2 Classification of Results

Based on the obtained literature (Sect. 4.1), we performed a qualitative coding in order to initial categorize each of the publications along the following dimensions: *sustainability* (is it an economical-, ecological- or social-oriented study?) [17], *level* (is it a

domain-specific or general approach?), *assessment* (is it quantitative or qualitative?), and *type of contribution* (what type of approach is proposed?) (based on [37]). To contribute to the robustness, the coding procedure was, as a first step, carried out by two researchers independently and, afterwards, consolidated in a follow-up workshop.

As indicated in Table 1, regarding *sustainability*, all articles address the economical dimensions (100%)—mainly by focusing on financial aspects—and the ecological (22%) and social (24%) characteristics are respected less. These findings underline our assumption that is stated in the Introduction, which says that most of the current articles focus on economic aspects while ecological and social ones are neglected. The financial emphasis is also visible regarding the kind of *assessment*, in which 34% are exclusively quantitative, 24% are qualitative and 40% used an approach that mixed these types. However, mixed approaches are often applied to quantify rather qualitative aspects, for instance [29] explores business alternatives and ranked each afterwards, or [50] identifies business model activities and translated them into measureable binary scheme.

Looking at the *level* dimension, most of the articles relate to a concrete business model instance (66%), for example analysing Smart Cities [16] or Walmart [40]. The minority provides general-purpose methods and criteria (36%) (e.g., [19, 57, 75]).

The *type of contribution* varies along different characteristics: framework (basic structure of a system/concept), construct (vocabulary/symbols like modelling grammar), model (representation of domain by using constructs), method (step-by-step guidance), instantiation (concrete implantations), and measure/criteria (value of a certain phenomenon) (based on [37]). Nevertheless, most of the studies develop or apply methods (58%) such as [15] who propose a 4-step approach to evaluate the consistency, as well as provide measures and indicators (56%), for example multidimensional evaluation criteria [47]. While 40% of the articles propose frameworks, 26% deal with mathematical models like DEA [75], 12% instantiations and 8% constructs.

4.3 Typology of Sustainability-Oriented Business Model Evaluation

Next, the gathered literature was analysed to identify approaches that facilitate sustainability-oriented business model evaluation. To structure our findings—a *total of 45 approaches*—, we used a typology (Fig. 2), which is usually derived in a deductive manner and gives a qualitative-based, conceptual classification [78]. Additionally, we visualized insights related to the *Frequency* (ratio and no. of studies using a type) and *Sustainability* (pillars that are addressed by a type). However, some studies combine different approaches and therefore aim to address more dimensions. For contributing to the robustness of the classification, we conducted a workshop with three researchers.

Most of the examined articles use a *comparison-oriented evaluation (I)* which can be a simple comparison of alternatives, a comparative analysis based on Boolean algebraic techniques—Qualitative Comparative Analysis—, a weighing up of trade-offs, or a rating of positive and negative impacts/consequences for business options. Additional techniques for solving multiple criteria decision-making problems are identified as *decision-structuring-oriented (VIII)* approaches such as Decision Matrix, Event-Decision Trees for problems with uncertainty and Morphological Box/Taxonomy.

Table 1. Overview of identified articles (● = characteristic fulfilled; – not fulfilled).

#Ref.	Sustainability		Level		Assessment			Type of contribution					Measure	
	Economic	Ecological	Social	Instance	General	Qualitative	Quantitative	Mixed	Framework	Construct	Model	Method		Instantiation
[38]	●	–	–	●	–	–	●	–	–	–	●	●	–	●
[39]	●	–	●	●	–	–	–	–	–	–	–	●	●	–
[18]	●	●	●	–	●	–	–	–	–	–	–	●	●	–
[40]	●	–	–	●	–	–	●	–	–	–	–	–	–	–
[41]	●	–	–	●	–	●	–	–	–	–	–	–	–	–
[42]	●	–	–	–	●	–	●	–	–	–	–	–	–	–
[43]	●	–	–	–	●	–	–	–	–	–	–	●	–	–
[44]	●	–	–	–	●	–	–	–	–	–	–	●	–	–
[16]	●	●	●	●	–	–	●	–	–	–	–	–	●	●
[45]	●	–	–	●	–	–	●	–	–	–	–	–	–	●
[46]	●	–	–	–	●	–	–	–	–	–	–	●	–	●
[31]	●	–	–	●	–	–	–	–	–	–	–	●	–	●
[47]	●	–	–	●	–	–	–	–	–	–	–	●	–	●
[48]	●	–	●	●	–	–	–	–	–	–	–	–	–	●
[6]	●	●	●	–	●	–	–	–	–	–	–	●	–	–
[33]	●	●	●	–	●	–	–	–	–	–	–	–	–	–
[49]	●	●	●	–	●	–	●	–	–	–	–	–	–	–
[50]	●	–	–	●	–	–	–	–	–	–	–	●	–	–
[51]	●	–	–	–	●	–	–	–	–	–	–	●	–	–
[52]	●	●	●	●	–	–	●	–	–	–	–	●	–	●
[29]	●	●	–	●	–	–	–	–	–	–	–	●	–	–
[53]	●	–	–	●	–	–	●	–	–	–	–	●	–	●
[54]	●	–	–	–	●	–	–	–	–	–	–	●	●	–
[55]	●	–	–	●	–	–	●	–	–	–	–	–	●	●
[56]	●	–	–	●	–	–	●	–	–	–	–	–	–	●

(continued)

Table 1. (continued)

#Ref.	Sustainability			Level		Assessment		Type of contribution				Measure		
	Economic	Ecological	Social	Instance	General	Qualitative	Quantitative	Mixed	Framework	Construct	Model		Method	Instantiation
[57]	●	-	-	-	●	-	-	●	●	-	-	-	-	●
[4]	●	●	●	-	●	-	-	●	●	-	-	-	-	●
[58]	●	-	-	●	-	-	-	●	●	-	-	-	-	-
[59]	●	-	-	-	●	-	-	●	●	-	-	-	-	●
[60]	●	-	-	-	●	-	-	●	●	-	●	-	-	●
[61]	●	-	-	-	●	-	-	●	●	-	-	-	-	●
[62]	●	-	-	-	●	-	-	●	●	-	-	-	-	-
[63]	●	-	-	●	-	●	-	●	●	-	-	●	-	●
[64]	●	●	-	●	-	●	-	●	●	-	-	-	-	-
[65]	●	●	-	●	-	-	-	●	●	-	-	-	-	-
[66]	●	-	-	●	-	-	-	-	●	-	●	-	-	●
[67]	●	-	●	●	●	-	-	●	●	-	●	-	●	●
[30]	●	-	-	●	-	-	-	●	●	-	●	-	-	-
[68]	●	-	-	●	-	-	-	●	●	-	-	-	-	●
[15]	●	-	-	●	-	-	-	●	●	-	-	-	-	●
[69]	●	-	-	●	-	-	-	●	●	-	-	-	-	-
[70]	●	-	-	●	-	-	-	●	●	-	-	-	-	●
[14]	●	-	-	●	-	-	-	-	●	-	-	-	-	-
[71]	●	-	-	●	-	-	-	●	●	-	-	-	-	●
[72]	●	-	-	●	-	-	-	●	●	-	-	-	-	●
[73]	●	-	-	●	-	-	-	-	-	-	-	-	-	●
[74]	●	-	-	●	-	-	-	-	-	-	●	-	-	●
[75]	●	●	●	-	●	-	-	-	-	-	●	-	-	●
[76]	●	-	-	●	-	-	-	-	-	-	●	-	-	●
[77]	●	-	●	●	-	-	-	-	-	-	●	-	-	●

Type/Description/References	Frequency		Sustainability		
	Ratio	No.	Econ.	Ecol.	Social
(I) Benchmark-, Comparison- and Trade Off-oriented Evaluation <i>Compare/Identify Alternatives</i> [6, 14, 15, 29, 30, 39, 40, 42–45, 47–49, 51, 61, 65, 66, 68, 70, 71, 74, 77], <i>Benchmark</i> [16], <i>Qualitative Comparative Analysis</i> [50], <i>Trade Off Analysis</i> [15, 65], <i>Collect/Value Impacts</i> [18, 43, 44, 47, 66]	64%	5	●	●	●
(II) Economic-/Financial-oriented Evaluation and Metrics <i>Return on Investment</i> [38, 55], <i>Life Cycle Costing</i> [71], <i>{Net Present Value</i> [55, 71], <i>Breakeven</i> [48], <i>Market Share</i> [48, 58, 74], <i>Valuation</i> [48], <i>Scalability</i> [74], <i>Channel Recurring</i> [74], <i>Transaction Sunk Cost</i> [74], <i>Profitability</i> [52], <i>Price-based Technical</i> [58], <i>Value-based Fundamental</i> [58], <i>Technology Competitiveness</i> [52], <i>Competitive Advantage</i> [74]; Analysis	36%	14	●	-	-
(III) Mathematical-oriented Evaluation Methods <i>Analytical Hierarchy Process</i> [29, 38, 51, 53, 54, 69, 72, 76, 77], <i>Fuzzy Comprehensive Evaluation Method</i> [51, 69, 72, 76], <i>DEA-model</i> [75], <i>Sensitivity Analysis</i> [38, 77]	32%	4	●	○	○
(IV) Survey- and Questionnaire-oriented Evaluation <i>Questionnaire</i> [16, 38, 45, 47, 52, 54, 59, 65], <i>Delphi Technique</i> [70, 76], <i>Group Decision</i> [6, 18, 69], <i>Consensual Assessment Technique</i> [45]	28%	4	●	●	●
(V) Simulation-based Evaluation Modelling Techniques/Tools <i>Simulation</i> [52, 54, 61, 66], <i>System Dynamics</i> [60], <i>Stress Test</i> [6], <i>Scenario Planning</i> [6, 30, 39, 45–47, 52, 70]	25%	4	●	-	-
(VI) Strategy-oriented Evaluation Tools <i>Balanced Scorecard</i> [19, 69], <i>PESTEL</i> [6], <i>PEST</i> [39, 52], <i>Five Forces Framework</i> [52], <i>SWOT Analysis</i> [57, 58, 62], <i>SUST-BMA</i> [19]	20%	6	●	●	●
(VII) Business Model Ontology-oriented Evaluation <i>Business Model Canvas</i> [6, 19, 62, 63, 67], <i>4C Net Framework</i> [73], <i>Ballon Framework</i> [65]	14%	3	●	-	-
(VIII) Decision Structuring-oriented Evaluation <i>Decision Matrix</i> [51, 61], <i>Event-Decision Tree</i> [31], <i>Morphological Box/Taxonomy</i> [29, 30, 47, 51]	14%	3	●	○	-
(IX) Pattern- and Key Question-based evaluation <i>Pattern-based Analysis</i> [46, 52], <i>Central/Key Questions</i> [59, 62]	8%	2	●	-	-
(X) Value Proposition-oriented Evaluation Tools <i>Value Mapping Tool</i> [18], <i>Value Proposition Traceability Diagram</i> [43, 44]	6%	2	●	●	●

Fig. 2. Resulting typology (Note: ● = fully fulfilled; ○ = partially fulfilled).

The *economic/financial-oriented evaluation (II)* is mainly based on metrics and calculations such as Return on Investment (ROI), Breakeven Analysis or Life Cycle Costing (LCC). *Mathematical-oriented evaluation (III)* makes use of methods like the Fuzzy Comprehensive Evaluation Method, Analytical Hierarchy Process (AHP) for generating priorities of decisions, and DEA-model for evaluating the effectiveness of decision-making. *Simulation-based (V)* approaches deal with System Dynamics Models, Scenario Planning, and Stress Tests to examine the robustness of an object.

Questionnaire and survey-oriented evaluation (IV) is designed for surveying individuals as well as for structuring group communication (e.g., Delphi Method or Amabile's Consensual Assessment Technique) aiming to achieve a group-based decision. *Pattern- and Key-question-based Evaluation (IX)* is used, for example to guide the analysis in respect of certain issues (e.g., profit patterns). Typically, such key questions base on business model components that are part of the *business model ontology (VII)* in which the Business Model Canvas or 4C Net Business Model Framework are applied, for example to specify certain analysis units.

Strategy-oriented (VI) tools focus mainly on the management level like the Balanced Scorecard that gives a balanced presentation of financial and operational

measures, and the SWOT Analysis that takes strengths, weaknesses, opportunities, and threats into account. In addition, strategic frameworks facilitate to analyze factors effecting an organization like PEST and PESTEL (macro-environment factors: political, economic, social, technological; for PESTEL additionally: environmental and legal) as well as the Five Forces Framework. Next to those strategy-related tools, there are also tools for analyzing the *value proposition (X)* and providing innovative forms of value.

5 Conclusion, Limitations and Research Agenda

Our contribution is a consolidated overview of the state-of-the art literature regarding sustainability-oriented business model evaluation in form of a typology. Although we derived helpful insights, our study is not free of limitations. Our investigation is limited to the outlets selected, which in our opinion does allow for a broad and interdisciplinary view. However, further sources may provide additional information. The identification and classification is based on own interpretations which have limitations—therefore, two researchers worked independently and consolidated afterwards.

Even though we focused on sustainability in a multidimensional manner, numerous articles gathered deal with economic aspects like costs, for which reason further research is needed. Therefore, we would argue that our work gives orientation as to which approaches and criteria can be applied or adapted to new fields. Our findings can be used, for example to (1) analyse the transferability of economical approaches to ecological and social aspects, (2) combine existing methods, (3) evaluate which approaches are easy to use and suitable for a domain, or (4) implement software- tools to facilitate continuous evaluation of business models. These efforts may lead to more holistic approaches that incorporate different sustainability dimensions. As a next step, we particular plan to evaluate selected approaches (e.g., group-based decisions) and extend our search to examine, if available, specific methods (e.g., in sustainability assessment).

Acknowledgments. This research was conducted in the scope of the research project “Smart-Hybrid—Process Engineering” (ZW 6-85003451), which is partly funded by the European Regional Development Fund and the State of Lower Saxony (NBank).

References

1. Brundtland, G.H.: World Commission on Environment and Development (WCED, 1987): Our Common Future. Oxford University Press, Oxford (1987)
2. Melville, N.P.: Information systems innovation for environmental sustainability. *MIS Q.* **34**, 1–21 (2010)
3. Seidel, S., Recker, J.C., vom Brocke, J.: Sensemaking and sustainable practicing: functional affordances of information systems in green transformations. *MIS Q.* **37**, 1275–1299 (2013)

4. Lüdeke-Freund, F., Freudenreich, B., Schaltegger, S., Saviuc, I., Stock, M.: Sustainability-oriented business model assessment—a conceptual foundation. In: Carayannis, E.G., Sindakis, S. (eds.) *Analytics, Innovation, and Excellence-Driven Enterprise Sustainability*, pp. 169–206. Palgrave Macmillan US, New York (2017)
5. Veit, D., Clemons, E., Benlian, A., Buxmann, P., Hess, T., Kundisch, D., Leimeister, J.M., Loos, P., Spann, M.: Business models: an information systems research agenda. *Bus. Inf. Syst. Eng.* **6**, 45–53 (2014)
6. Haaker, T., Bouwman, H., Janssen, W., de Reuver, M.: Business model stress testing: a practical approach to test the robustness of a business model. *Futures* **89**, 14–25 (2017)
7. Schaltegger, S., Hansen, E.G., Lüdeke-Freund, F.: Business models for sustainability: origins, present research, and future avenues. *Organ. Environ.* **29**, 3–10 (2016)
8. Osterwalder, A., Pigneur, Y.: *Business Model Generation: A Handbook for Visionaries, Game Changers, and Challengers*. Wiley, Hoboken (2010)
9. Schoormann, T., Behrens, D., Kolek, E., Knackstedt, R.: Sustainability in business models—a literature-review-based design-science-oriented research agenda. In: *Proceedings of the European Conference on Information Systems (ECIS)*, Istanbul, Turkey (2016)
10. Täuscher, K., Abdelkafi, N.: Visual tools for business model innovation: recommendations from a cognitive perspective. *Creat. Innov. Manag.* **26**, 160–174 (2017)
11. Afuah, A., Tucci, C.L.: *Internet Business Models and Strategies*. McGraw-Hill, New York (2001)
12. Kamoun, F.: Rethinking the business model with RFID. *Commun. Assoc. Inf. Syst.* **22**, 35 (2008)
13. Gordijn, J., Akkermans, H.: Designing and evaluating e-business models. *IEEE Intell. Syst.* **16**, 11–17 (2001)
14. Tesch, J., Brillinger, A.: The evaluation aspect of the digital business model innovation: a literature review on tools and methodologies. In: *Proceedings of the European Conference on Information Systems*, pp. 2250–2268 (2017)
15. Si, Y., Chen, T., Zheng, W.: An approach to evaluating business models - a case study of Taobao. In: *2011 IEEE 18th International Conference on Industrial Engineering and Engineering Management*, pp. 31–35 (2011)
16. Díaz-Díaz, R., Muñoz, L., Pérez-González, D.: The business model evaluation tool for smart cities: application to SmartSantander use cases. *Energies* **10**, 262 (2017)
17. Isaksson, R., Garvare, R.: Measuring sustainable development using process models. *Manag. Audit. J.* **18**, 649–656 (2003)
18. Bocken, N.M.P., Rana, P., Short, S.W.: Value mapping for sustainable business thinking. *J. Ind. Prod. Eng.* **32**, 67–81 (2015)
19. Lüdeke-Freund, F., Freudenreich, B., Schaltegger, S., Saviuc, I., Stock, M.: Sustainability-oriented business model assessment—a conceptual foundation. In: *Analytics, Innovation, and Excellence-Driven Enterprise Sustainability*, pp. 169–206. Palgrave Macmillan, New York (2017)
20. Hamel, G.: *Leading the Revolution*. Harvard Business School Press, Boston (2000)
21. Wirtz, B.W.: *Business Model Management, Design-Instrumente-Erfolgsfaktoren von Geschäftsmodellen* (2011)
22. Zott, C., Amit, R.: Measuring the Performance Implications of Business Model Design: Evidence from Emerging Growth Public Firms. Insead, Fontainebleau (2002)
23. Alexa, M.: Business model evaluation—a conceptual approach. *Rev. Econ. Bus. Stud.* **14**, 245–260 (2014)
24. Magretta, J.: *Why business models matter* (2002)
25. Osterwalder, A., Pigneur, Y., Tucci, C.L.: Clarifying business models: origins, present, and future of the concept. *Commun. Assoc. Inf. Syst.* **16**, 1 (2005)

26. Lüdeke-Freund, F., Dembek, K.: Sustainable business model research and practice: emerging field or passing fancy? *J. Clean. Prod.* **168**, 1668–1678 (2017)
27. Stubbs, W., Cocklin, C.: Conceptualizing a “sustainability business model”. *Organ. Environ.* **21**, 103–127 (2008)
28. Upward, A., Jones, P.: An ontology for strongly sustainable business models: defining an enterprise framework compatible with natural and social science. *Organ. Environ.* **29**, 97–123 (2016)
29. Jun, M., Di Muro, A.: Holistic methodology to analyze EV business models. *Int. J. Innov. Manag. Technol.* **4**, 259 (2013)
30. Seim, S., Scheller, F., Götz, M., Kondziella, H., Bruckner, T.: Assessment of PV-based business models in urban energy systems with respect to political and economic targets: a model-based scenario analysis. In: 10 Internationale Energiewirtschaftstagung IEWT 2017 Wien (2017)
31. Copani, G., Rosa, P.: DEMAT: sustainability assessment of new flexibility-oriented business models in the machine tools industry. *Int. J. Comput. Integr. Manuf.* **28**, 408–417 (2015)
32. Weill, P., Vitale, M.: *Place to Space: Migrating to eBusiness Models*. Harvard Business Press, Boston (2001)
33. Hansen, E.G., Grosse-Dunker, F., Reichwald, R.: Sustainability innovation cube—a framework to evaluate sustainability-oriented innovations. *Int. J. Innov. Manag.* **13**, 683–713 (2009)
34. Ness, B., Urbel-Piirsalu, E., Anderberg, S., Olsson, L.: Categorising tools for sustainability assessment. *Ecol. Econ.* **60**, 498–508 (2007)
35. Webster, J., Watson, R.T.: Analyzing the past to prepare for the future: writing a literature review. *MIS Q.*, xiii–xxiii (2002)
36. vom Brocke, J., Simons, A., Niehaves, B., Riemer, K., Plattfaut, R., Cleven, A.: Reconstructing the giant: on the importance of rigour in documenting the literature search process. In: *Proceedings of the European Conference on Information Systems*, Verona, Italy, pp. 2206–2217 (2009)
37. Hevner, A., March, S., Park, J., Ram, S.: Design science in information systems research. *MIS Q.* **28**, 75–105 (2004)
38. Ali, A.: An MCDM approach towards m-payment business models evaluation. *Int. J. Anal. Hierarchy Process.* **7**, 273–294 (2015)
39. Bilgeri, D., Brandt, V., Lang, M., Tesch, J., Weinberger, M.: *The IoT business model builder*. A White Paper of the Bosch IoT Lab in collaboration with Bosch Software Innovations GmbH (2015)
40. Brea-Solis, H., Casadesus-Masanell, R., Grifell-Tatjé, E.: Business model evaluation: quantifying Walmart’s sources of advantage. *Strateg. Entrep. J.* **9**, 12–33 (2015)
41. Cagliano, A.C., De Marco, A., Mangano, G., Zenezini, G.: Assessing City Logistics projects: a business-oriented approach. In: *Proceedings of the Summer School “Francesco Turco”*, pp. 204–208. Italian Association of Industrial Operations Professors, Napoli (2016)
42. Clauss, T.: Measuring business model innovation: conceptualization, scale development, and proof of performance. *RD Manag.* **47**, 385–403 (2017)
43. Costa, C., Cunha, P.: The social dimension of business models: an Actor-Network Theory perspective. In: *Proceedings of the American Conference on Information Systems (AMCIS)* (2015)
44. Costa, C.C., Cunha, P.R.: More than a gut feeling: Ensuring your inter-organizational business model works. In: *Proceedings of the Bled eConference*, p. 31 (2015)
45. Ebel, P., Bretschneider, U., Leimeister, J.M.: Leveraging virtual business model innovation: a framework for designing business model development tools: leveraging virtual business model innovation. *Inf. Syst. J.* **26**, 519–550 (2016)

46. Echterfeld, J., Amshoff, B., Gausemeier, J.: How to use business model patterns for exploiting disruptive technologies. In: 24th International Conference on Management of Technology, International Association for Management of Technology, pp. 2294–2313 (2015)
47. Ghezzi, A., Georgiades, M., Reichl, P., Le-Sauze, N., Cairano-Gilfedder, C.D., Managiaracina, R.: Generating innovative interconnection business models for the future internet. *info* **15**, 43–68 (2013)
48. Grustam, A.S., Vrijhoef, H.J.M., Koymans, R., Hukal, P., Severens, J.L.: Assessment of a Business-to-Consumer (B2C) model for Telemonitoring patients with Chronic Heart Failure (CHF). *BMC Med. Inf. Decis. Mak.* **17**, 145 (2017)
49. Hart, S.L., Milstein, M.B.: Creating sustainable value. *Acad. Manag. Exec.* **17**, 56–67 (2003)
50. Hvass, K.: A Boolean Approach to Airline Business Model Innovation. Copenhagen: CIBEM Working Paper Series 13 (2012)
51. Im, K., Cho, H.: A systematic approach for developing a new business model using morphological analysis and integrated fuzzy approach. *Expert Syst. Appl.* **40**, 4463–4477 (2013)
52. Ishida, F., Sakuma, H., Abe, H., Fazekas, B.: Remodeling method for business models of R D outputs. In: 2006 Technology Management for the Global Future - PICMET 2006 Conference, pp. 708–714 (2006)
53. Kampker, A., Heimes, H.H., Ordnung, M., Lienemann, C., Hollah, A., Sarovic, N.: Evaluation of a remanufacturing for lithium ion batteries from electric cars. *World Acad. Sci. Eng. Technol. Int. J. Mech. Mechatron. Eng.* **3** (2016)
54. Kayaoglu, N.: A Generic Approach for Dynamic Business Model Evaluation (2013)
55. Laplume, A., Anzalone, G.C., Pearce, J.M.: Open-source, self-replicating 3-D printer factory for small-business manufacturing. *Int. J. Adv. Manuf. Technol.* **85**, 633–642 (2016)
56. Lee, Y.H., Shim, H.D., Park, N.I.: A new convergence business model framework for context-aware computing service market. *Commun. Converg. Rev.* **2**, 1–12 (2010)
57. Leitão, A., Cunha, P., Valente, F., Marques, P.: Roadmap for business models definition in manufacturing companies. *Procedia CIRP* **7**, 383–388 (2013)
58. Martínez, A.B., Galván, R.S., Alam, S.: Financial analysis of retail business organization: a case of Wal-Mart Stores, Inc. *Nile J Bus. Econ.* **3**, 67–89 (2017)
59. Mateu, J.M., March-Chorda, I.: Searching for better business models assessment methods. *Manag. Decis.* **54**, 2433–2446 (2016)
60. Moellers, T., Bansemir, B., Pretzl, M., Gassmann, O.: Design and evaluation of a system dynamics based business model evaluation method. In: Maedche, A., vom Brocke, J., Heyvner, A. (eds.) *DESIRIST 2017. LNCS*, vol. 10243, pp. 125–144. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-59144-5_8
61. O Riordan, N., O'Reilly, P., Duane, A., Andreev, P.: Business model innovation: a temporal perspective. In: Proceedings of the Australasian Conference on Information Systems, Auckland New Zealand, 8–10 December 2014 (2014)
62. Osterwalder, A.: How to describe and improve your business model to compete better. *Trobe Univ. Acesso Em.* **26**, 2015 (2007)
63. Prado, A.M., Calderon, D., Zúñiga, R.: Providing low-cost and high-quality medications to rural communities in developing countries: the case of Accion Medica Cristiana in Nicaragua. *J. Bus. Res.* **69**, 3910–3922 (2016)
64. Qin, Q., Liang, F., Li, L., Wei, Y.-M.: Selection of energy performance contracting business models: a behavioral decision-making approach. *Renew. Sustain. Energy Rev.* **72**, 422–433 (2017)

65. Raju, A., Lindmark, S., Yaron, O., De Poorter, E., Tytgat, L., Delaere, S., Ballon, P.: Business model assessment of Green wireless sensor ecosystems. In: 11th Conference of Telecommunication, Media and Internet Techno-Economics (CTTE-2012), pp. 1–8. Ghent University, Department of Information Technology (2012)
66. Sánchez, P., Ricart, J.E.: Business model innovation and sources of value creation in low-income markets. *Eur. Manag. Rev.* **7**, 138–154 (2010)
67. Schüle, S., Schubert, M., Hoyer, C., Dressel, K.-M.: Development of an assessment tool to evaluate and improve SME business models. *J. Bus. Models* **4**, 5 (2016)
68. Sharma, S., Gutiérrez, J.A.: An evaluation framework for viable business models for m-commerce in the information technology sector. *Electron. Mark.* **20**, 33–52 (2010)
69. Su, C.H., Hung, Y.H., Tzeng, G.H.: Fuzzy multiple attribute decision making theory with the balanced scorecard application in mobile industry. In: 2011 IEEE International Conference on Fuzzy Systems (FUZZ-IEEE 2011), pp. 1479–1484 (2011)
70. Suikki, R.M., Goman, A.M.J., Haapasalo, H.J.O.: A framework for creating business models – a challenge in convergence of high clock speed industry. *Int. J. Bus. Environ.* **1**, 211–233 (2006)
71. Van Ostaeyen, J., Dufloy, J.: Assessing the potential of business model innovation for investment goods through Life Cycle Costing. In: Proceedings of the CIRP IPS2 Conference, pp. 527–534. Linköping University Electronic Press (2012)
72. Wang, X., Brooks, S., Sarker, S.: A review of green IS research and directions for future studies. *Commun. Assoc. Inf. Syst.* **37**, 395–429 (2015)
73. Wirtz, B.W., Lihotzky, N.: Customer retention management in the B2C electronic business. *Long Range Plann.* **36**, 517–532 (2003)
74. Young, L.W., Johnston, R.B.: A framework for evaluating the impact of the internet on business-to-business e-commerce value delivery options. *J. Syst. Inf. Technol.* **7**, 129–152 (2003)
75. Zeng, D.-B., Lu, H.-X.: Evaluation Study of Internet Commercial Sustainable Development under the New Commercial Civilization (2014)
76. Zhang, G., Jiang, Y., Chen, X., Wang, Z.: The research on business model evaluation based on internet free services. In: 2012 Proceedings of Computer Science and Information Management (ICIIM 2012), pp. 97–102 (2012)
77. Zografos, K.G., Androutsopoulos, K.N., Sihvola, T.: A methodological approach for developing and assessing business models for flexible transport systems. *Transportation* **35**, 777–795 (2008)
78. Bailey, K.D.: *Typologies and Taxonomies: An Introduction to Classification Techniques*. Sage, Thousand Oaks (1994)



Towards Agility in IT Governance Frameworks

Sulejman Vejseli¹ and Alexander Rossmann²

¹ University of the West of Scotland, Paisley Campus, Paisley PA1 2BE, Scotland, UK
sulejman.vejseli@hhz.de

² Reutlingen University, Alteburgstrasse 150, 72762 Reutlingen, Germany
Alexander.Rossmann@Reutlingen-University.de

Abstract. Digital transformation has changed corporate reality and, with that, firms' IT environments and IT governance (ITG). As such, the perspective of ITG has shifted from the design of a relatively stable, closed and controllable system of a self-sufficient enterprise to a relatively fluid, open, agile and transformational system of networked co-adaptive entities. Related to this paradigm shift in ITG, this paper aims to clarify how the concept of an effective ITG framework has changed in terms of the demand for agility in organizations. Thus, this study conducted 33 qualitative interviews with executives and senior managers from the banking industry in Germany, Switzerland and Austria. Analysis of the interviews focused on the formation of categories and the assignment of individual text parts (codings) to these categories to allow for a quantitative evaluation of the codings per category. Regarding traditional and agile ITG dimensions, 22 traditional and 25 agile dimensions were identified. Moreover, agile strategies within the agile ITG construct and ten ITG patterns were identified from the interview data. The data show relevant perspectives on the implementation of traditional and new ITG dimensions and highlight ambidextrous aspects in ITG frameworks.

Keywords: Agility · IT governance · Ambidexterity · Agile strategies

1 Introduction

Digital transformation has fundamentally changed corporate reality. The role of information technology (IT) in business is now pervasive. Whereas the dominant logic of the Industrial Age was linear and product-oriented, the Digital Age is nonlinear and service-oriented [1]. This broad paradigm shift in logic is also reflected in changing IT environments and, consequently, in firms' IT governance (ITG). The perspective has shifted from the relatively stable, closed and controllable system of a self-sufficient enterprise to the relatively fluid, open, agile and transformational system of networked co-adaptive entities [2]. Thus, the demand for agility has become more important to enterprises. Existing studies have recognised efficiency and stability as core concepts in ITG design [3]. This makes sense as the “old” world was characterized by a stable, placid environment, in which neither the core technology nor the markets in which companies were operating changed drastically over time. Organizations could afford to use “command-and-control” mechanisms to govern IT [3]. In this context, many proposed methodologies, reference guides,

sets of best practices (e.g. COBIT), and frameworks such as the IT Infrastructure Library (ITIL) have emerged and developed in recent years. However, the adoption of such ITG models, or so-called conventional or traditional frameworks, does not necessary yield the desired return in the context of digital transformation [4]. By contrast, agile processes and methods have evolved in the past years, especially in the area of software development [5, 6]. Independent of the business area, these agile frameworks can “add value” to business organizations, through a process in which the principles of communication and collaboration are essential [7]. Thus, adopting the agile principles, values and best practices to the context of ITG can bring even more meaningful results to organizational management. Their benefits can lead to an increase in the speed of decision making, the insurance of business processes, organizational competitiveness and other aspects [4]. Thus, agile ITG has become highly relevant to keep up with competitors in today’s dynamic world. To date, the broad scientific community has not analysed the impetus of agile strategies on ITG [4, 8]. However, research interest in this topic is growing [5]. As such, this study aims to investigate the impact of agility on effective ITG for today’s dynamic world, which yields the following research question (RQ): *How is the concept of an effective ITG framework changing in response to the demand for agility in organizations?*

To answer this RQ, we analysed several traditional and agile aspects of governance dimensions gleaned from 33 qualitative interviews with executives and senior managers from the banking industry in Germany, Switzerland and Austria. In doing so, we were able to elicit major patterns for ITG dimensions in the digital world. Strategies used to implement an agile ITG can be grouped under the same general dimensions as in the conventional ITG literature, which enables us to compare traditional and agile aspects and to derive patterns for the effectiveness of an agile ITG framework. Such patterns help explain “real-world” problems because they capture and allow for reuse of experiences of best practices in a specific professional domain [9]. Thus, the contribution of this paper is threefold. First, the study uncovers agile ITG dimensions that complement the traditional elements of ITG; second, it provides an overview of agile strategies used within the governance construct; third, it elicits patterns for effective ITG dimensions.

2 Theoretical Background

2.1 ITG and Its Dimensions

Since late 1990s, many researchers and practitioners have investigated ITG from different perspectives [10]. As such, several definitions have been proposed. While some definitions such as those from Peterson [3] and Weill and Ross [11] focus on the decision-making process within the ITG framework and do not address any role aspects, other definitions such as those from van Grembergen and de Haes [12] and the IT Governance Institute [13] adequately address objectives, objects and subjects of ITG decisions [14]. We argue that an ITG definition should include both structure and process aspects. Therefore, enhancing the definition of the IT Governance Institute with Weill and Ross’s characterization should help cover the most relevant dimensional concepts of current ITG research, leading to the following definition:

ITG is the responsibility of executives and the board of directors and consists of the leadership and organizational structures and processes that ensure that the organization's IT sustains and extends firm strategies and objectives. ITG represents the framework for decision rights and accountabilities to encourage desirable behaviour in the use of IT.

With this definition, implementing ITG effectively requires a set of ITG instruments to gain congruence with the firm's mission, strategy, values, norms and culture [15], which in turn leads to desirable IT behaviours and governance outcomes [11]. However, implementing ITG is a complex issue because it is contingent on a variety of sometimes conflicting internal and external factors [16–18]. Effective ITG doesn't happen by accident and top-performing firms should consequently carefully design governance [17]. The effectiveness of a firm's ITG can be assessed by evaluating how well it enables IT to deliver on four objectives: cost-effectiveness, asset utilization, business growth and business flexibility [17]. Consequently, in the literature several studies have argued that organizations should use ITG [15, 19, 20]. Therefore, consistent with the previous definition and the literature, implementing an effective ITG requires a framework based on three major dimensions [21].

Structure. The framework needs to answer the following questions: Who makes the decisions? Which organizational units will be created? Who will take part in these organizational units? What responsibilities will they assume [22]? Examples of traditional structures are IT steering committees, IT project steering committees and IT strategy committees as well as structures that enable CIOs to report to CEOs [15].

Process. The process aspect targets the following questions: How are IT investment decisions made? What are the decision-making processes for proposing, reviewing, approving and prioritizing investments? Conventional processes, for example, contain portfolio management, IT budget control and reporting, project governance methodologies or information systems planning [23].

Communication/Relational Mechanisms. The aspects dealing with communication and relational mechanisms pose the question of how the results of ITG processes and decisions will be monitored, measured and communicated. Also required are mechanisms to communicate IT investment decisions to the board of directors, executive management, business management, IT management, employees and shareholders [24]. Examples of traditional communication/relational mechanisms are a shared understanding of business/IT objectives, cross-functional business/IT training and collaboration between principal stakeholders [3].

Thus, deploying ITG in a firm means using a mixture of various structures, processes and relational mechanisms. Therefore, in the past decades, several frameworks that support the implementation of ITG have been created. Some of the most familiar frameworks are COBIT and ITIL. However, no comprehensive framework covers all structures, processes and relational mechanisms for a comprehensive ITG approach [25]. Depending on a firm's structure, a mixture of different elements is required [19]. Furthermore, in terms of flexibility agile strategies can play a significant role in designing ITG.

2.2 Ambidexterity of ITG

The ability to align governance structures to existing capabilities and environmental conditions is an important success factor for digital transformation, and the alignment procedures need to be agile. As Gersick [26], Romanelli and Tushman [27] and Greiner [28] showed, successful organizations alternate between two states in their organizational development. The first state is characterized by a phase of environmental stability, in which firms strive for optimization within their existing business logic by focusing on minor, incremental adjustments in operational efficiency and benefit from economies of scale [26, 29]. The second state is characterized by fast-changing and highly volatile environmental conditions, in which major organizational adjustment are required to successfully manage the change [29]. This view of organizational development is referred to as the punctuated equilibrium theory. Depending on which state the organization is in, it should design adequate governance dimensions to deal successfully with the specific challenges of each state [30]. Thus, the distinction of these two states is important for the design of an appropriate governance structure and the execution of measures and mechanisms [26]. Adequate governance structures are a crucial factor for company success, especially when the economy, society, technology and regulations undergo fundamental and highly complex changes [31], and are also dependent on the status of the external environment and internal capabilities [26].

One way for organizations to deal with these two types of states is by becoming an ambidextrous organization [29]; here, organizations have the ability “to both explore and exploit – to compete in mature technologies and markets where efficiency, control, and incremental improvement are prized and to also compete in new technologies and markets where flexibility, autonomy, and experimentation are needed” [32]. During the exploitation state, the organization focuses on activities to improve efficiency and reduce variance, while it concentrates on discovery and innovation activities in the exploration state. Especially incumbent organizations can benefit from taking an ambidextrous approach to organizational development [33]. The challenge of becoming ambidextrous lies in the organization’s ability to balance the two opposing state characteristics in existing business activities to sufficiently monetise them, while exploring new market opportunities to stay competitive in the future [34]. Tushman and O’Reilly [29] define organizational ambidexterity as a concept of structurally separated business units with distinct organizational responsibility. However, Markides [33] and Markides and Charitou [35] criticise this view for not considering the positive impact of potential benefits from synergies within business units in its assumption of strict separation. Consequently, researchers have introduced ambidexterity models with less strict separation. This contextual dimension of ambidextrous organizations is also manifested in the dual operating system formulated by Kotter [36]. Traditional hierarchical governance structures alone are not able to cope with the increase in speed of change. Thus, traditional hierarchical governance structures should be complemented by network-like governance structures that can react quickly to changes in the organization’s environment. The following sections follow the mindset of the punctuated equilibrium theory by referring to agile ITG as traditional ITG dimensions complemented by agile approaches.

2.3 Towards an Agile ITG

Many researchers have argued that business agility is required to survive the voracity of the market [37, 38]. Agility is important to change the direction of the environment and respond efficiently and effectively to such changes [39]. Consequently, in recent years the term “agile” has gained much attention from practitioners and academics because of its importance to the innovation and competitive performance of firms in contemporary business environments [40]. As business agility is a complex, multidimensional and context-specific concept, the literature has proposed several different concepts, frameworks and metrics for de-fining and explaining it [41]. Some empirical literature argues that enterprise agility is a kind of dynamic capability that enables a firm to reconfigure, assemble and integrate resources, information, processes and technologies that are embedded in different activities within an enterprise or its subsidiaries [42]. Using this ability enables a firm to create synergies and additional competitive advantages that enhance firm performance [40]. Other research argues that agility involves a firm’s capabilities related to interactions with customers, orchestration of internal operations and utilization of its ecosystem of external business partners [40, 43].

However, limited research has combined agile capabilities with governance frameworks. Qumer [6] introduced a first definition of agile governance, focusing on agile software development. He presents a summary of exploratory reviews and analysis to identify the related concepts, key aspects and importance of ITG. A second definition was presented by Cheng et al. [5], who introduced a list of measurement and control aspects in agile governance frameworks and validated this view with three case studies. The third definition of agile governance was proposed by Luna et al. [4], who specifically focused on ITG. In another study, Luna et al. [44] presented a fourth definition on a holistically and widely agile governance. The study of Luna et al. [45] offers an overview of nascent research with respect to agile strategies with-in governance and is grounded in a systematic literature review. We adopt the approach of Luna et al. [4] on the agile strategies in ITG because it specifically focuses on ITG and is the most comprehensive framework provided in the literature. In their research, Luna et al. [4] present a new concept of agile ITG in which principles, values and practices of the agile paradigm from software engineering are translated into the context of ITG. They suggest enhancing the ITG dimensions by the values and principles of the agile manifesto of software engineering introduced by Beck et al. [46]. They further argue that it is enough to adjust the focus of existing traditional practices, such as COBIT and ITIL, to agree with the principles and values supported by the agile manifesto and with the application of best practices that can be adapted from agile software engineering. Madi et al. [47] provide a list of the extracted values from the agile manifesto. This list is based on the foundation of agile practices and principles and contains the following values: “(1) Collaboration, (2) Communication, (3) Working software, (4) Flexibility, (5) Customer-centric, (6) Incremental, (7) Iterative, (8) Motivation, (9) Respect, (10) Trust, (11) Feedback, (12) Speed, (13) Technical excellence, (14) Simplicity, (15) Self-organizing, and (16) Learning”.

The manifesto and all its values and principles represent the philosophy behind agile strategies and ideally should be present in all practices proposed by various agile methods [48].

Therefore, the core objective of the current study is to try to obtain traditional as well as agile strategies put in place in firms, to identify how the concept of an effective framework is changing with respect to the demand for agility in organizations. Then, to elicit ITG patterns, the traditional dimensions and agile strategies must be set in context to some control variables. This is important because determining the right ITG dimensions is a complex endeavour, and it should be recognised that what strategically works for one company does not necessarily work for another [19]. Therefore, to elicit some ITG patterns of success, this study uses control variables such as firm size, IT strategy and regional differences in context, as also highlighted in the approaches of Weill [20] and Sambamurthy and Zmud [16].

3 Research Method

This research follows a qualitative approach in which we conducted semi-structured interviews with top management executives in the German-speaking banking industry. According to Benbasat [49], a qualitative research approach is useful for addressing the “how” question in the exploratory stage of knowledge building. This is particularly applicable to the exploration of (1) relevant governance constructs and (2) differences between traditional and agile governance dimensions from an empirical standpoint. Therefore, according to the formulated RQ, this work can be classified as exploratory research and interviews with experts are a well-established method in this context [50].

The first stage of research entails the development of a semi-structured questionnaire to provide an interview guideline. Workshops with researchers from different research institutions were conducted to formulate a first draft of the questionnaire. To ensure the suitability of the questions, five pre-interviews with top management executives in German-speaking banks were conducted by telephone. All interviews were recorded and fully transcribed. The discussion of the transcripts during the first stage led to adjustments according to the output of this stage.

In the second stage, a database with banks in Germany, Switzerland and Austria was prepared in order to send a letter of solicitation to the targeted industry. On the one hand, the banking industry is appropriate for research on ITG in digital transformation because many banks are currently redesigning their governance frameworks. On the other hand, business models in the banking industry are not fundamentally different from bank to bank. Therefore, a focus on this industry allows for an investigation of a well-defined context under comparable conditions within the whole sample. More than 1,800 e-mails were sent out, leading to initial contacts with more than 50 banks. Furthermore, it was important to identify the top executive of each bank mainly responsible for digital transformation to be included in the interview procedure. This step resulted in a final sample size of 33 executive interviews with banks in Germany, Switzerland and Austria.

Accordingly, in stage three, the 33 interviews were conducted by telephone. Before each interview, the questionnaire was given to the interviewee to serve as a guide.

Furthermore, all respondents gave permission to have the interview digitally recorded. Thus, all 33 semi-structured interviews were conducted and audio-recorded by two members of the panel. Voice recording is essential because it saves the recorded information at hand, thus ensuring accurate information of the data [51]. Each interview session took approximately 30 min. All the interviewees were at managerial levels and held executive or senior managerial positions with decision rights to define digital transformation projects.

In stage four, the audio records were fully transcribed to enable analysis of the data set. The transcribed source texts were first analysed case by case, to set the basis for the development of a system of categories for a structured evaluation of the text material. In doing so, the overall approach corresponds to the approach of a qualitative data analysis [52]. From a methodological point of view, the qualitative data analysis is based on the formation of categories and the assignment of individual text parts to these categories. To support the encoding of the data, analysis software MAX QDA was used. The transcribed texts were first imported into MAX QDA, and then individual text parts (= codings) of the transcribed interviews were assigned to the defined categories. When necessary, the main categories were further differentiated into subcategories, so that the data could be organized and interpreted in a category-based manner. In doing so, the quantitative evaluation referred to the quantity of the codings per category. Consequently, the data per category were quantitatively evaluated and qualitatively interpreted. The thematic codings allowed us to highlight the links and concepts related to the organizational integration of governance dimensions in the Digital Age. Many dimensions, especially the traditional ITG dimensions, were confirmed by the literature mentioned previously. However, our analysis led to several new elements of governance, mainly in the context of agility.

4 Data Analysis

Data analysis uncovered several traditional and agile ITG dimensions mentioned by the respondents. Table 1 provides an overview of the bank characteristics and control variables. We included firm size, degree of agility, duration of strategic focus on digital transformation, and regional differences as control variables in our questionnaire. Different categories of banks are well distributed on all control variables, with a slight focus on small and medium-sized banks (<500 employees). A differentiation between subsamples and categories per control variable extends the scope of this paper but opens a fruitful avenue for further analysis.

Table 2 lists all the information gathered from the 33 banks regarding different ITG dimensions. Overall, 47 dimensions for ITG appear in the sample. The ITG dimensions used by the firms are grouped into the categories “structures”, “processes” and “relational mechanisms”. To group the dimensions into the three mentioned categories, the study of Almeida et al. [23] served as a template. Many other studies, such as that of de Haes and van Grembergen [53] and Symons [21], use this type of framework to structure ITG dimensions. Moreover, we subdivided the dimensions into “traditional” and “agile” to differentiate between these two subcategories. To identify traditional ITG dimension, we used the framework of Almeida et al. [23].

Table 1. Bank characteristics and control variables

Control Variables		Banks																																		
		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33		
Firm size	Small (<500 employees)	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	
	Medium-sized (500–2000 employees)	•			•		•		•													•											•	•	•	
	Large (>2000 employees)						•					•	•									•	•			•										
Degree of agility in ITG	Very agile	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	
	A little agile	•			•																															
	Not agile				•							•																								
IT Strategy focus on digitalisation since	<3 years	•	•							•			•									•														
	3–4 years	•		•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	
	>4 years																																			
Regional differences	Germany	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	
	Switzerland	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
	Austria																																			

Regarding the agile ITG dimensions, we found no corresponding template in the literature. Therefore, a section in the questionnaire specifically asked the respondents about the agile elements within the ITG dimensions to identify items in this domain. This allowed us to recognise if an uncovered dimension was agile or not.

We set the following specifications: If a dimension does not exist, the cell is empty; when the dimension is implemented or there is some evidence that it is used, the cell is marked with a symbol “•”.

4.1 Traditional ITG Dimensions

The interviewees reported on how the goals were inherent in their business strategy and how they strived to manage digital transformation projects. Therefore, in exploring the traditional ITG dimensions implemented by banks in Germany, Switzerland and Austria we identified 22 traditional ITG dimensions (six structure elements, seven processes and nine relational/communication mechanisms; see Table 2). Many of the elements are also confirmed in the literature [23].

In terms of structure, the most frequently mentioned element was “CIO/COO on executive committee”. Of the 33 interviewees, 26 noted that the CIO or COO is a full member of the executive committee or has a direct reporting structure to the CEO. This situation allows a firm to ensure that IT is a regular agenda item and reporting issue for the board of directors. Furthermore, boards and committees (25 codings) are formed to help align business and IT and determine business priorities in IT investments. To prioritise and manage IT projects, project steering committees (18 banks) and IT organization structures (17 banks) are primarily used. Surprisingly, only in few institutions the CIO is a member of the executive board (3 codings).

Regarding traditional processes, portfolio management was highlighted as the most essential element. Of the 33 interviewees, 32 mentioned that within portfolio management, in which business and IT are involved, IT investments and projects are prioritised. In terms of controlling, monitoring and reporting projects, most banks use the processes “IT budget control and reporting” (20 codings), “project governance” (20 codings) and

“project tracking” (15 codings). Formal processes to define and update IT strategy are implemented in almost half the institutions (process: strategic information systems planning; 16 codings). In relatively few banks are the processes of benefit management (8 codings) and demand management (6 codings) implemented.

Communication in the financial institutions is crucial. Of the 33 interviewees, 27 highlighted the importance of regular internal communication. Here, leadership aspects play a key role. Seventeen interviewees indicated that management should give good examples in communicating. Furthermore, 14 executives mentioned IT leadership as important to articulate a vision for IT’s role in the company and to ensure that this vision is clearly understood by managers throughout the organization.

To ensure that business and IT work together, more than 50% of the firms locate business and IT close to each other (business/IT collocation; 17 codings). Only seven banks use cross-functional business/IT training instead. Ten banks use informal meetings of executives and senior management announcements. Finally, seven interviewees mentioned practicing knowledge management in their firms.

4.2 Agile ITG Dimensions

This subsection presents the analysis of the bank executives’ perceptions of the agile ITG dimensions. From their answers to the questionnaire, we identified six structure elements, ten process dimensions and nine relational/communication mechanisms.

Regarding structures, executives from 13 banks each mentioned the dimensions “digital transformation units”, “short and flexible decision paths” and “interdisciplinary/small project teams”. Similarly important, executives from 11 banks mentioned the dimension “lean project structures”. The respondents noted that the setup of new dedicated units for digital transformation allows for better communication and more intensive collaboration. Relatively few institutions use matrix structures (3 codings). In general, such units are created from organizational structures and include various positions from other units. However, short and flexible decision paths are important for speediness and flexibility in several processes. Such aspects are also reflected in the dimensions “interdisciplinary/small project teams” and “lean project structures”. Overall, structure elements should be kept simple to allow agile decision making.

In terms of processes, the respondents noted to use agile practices (20 codings). Many bank executives indicated that they used agile strategies such as from scrum, design thinking, lean approach, and so on. The executives further noted that taking higher risks by following trial-and-error processes (17 codings) might enhance the company’s agility. Such processes stimulate people to act in self-organized ways and ensure continuous learning. Furthermore, respondents highlighted fast and agile decision-making processes (15 codings) as significant to act in a flexible and speedy manner when making decisions. Moreover, 10 banks use innovative key performance indicators (KPIs), such as conversion rates or online client feedback, to improve their processes. Equally important were the introduction of lessons learned and innovation processes (10 codings). Few executives mentioned using more flexible and faster prioritizing, coordination or project monitoring and evaluation processes.

One of the most relevant dimensions, in terms of relational mechanisms, the respondents highlighted was external collaboration. Seventy-six percent of the executives agreed that collaboration (25 codings) with external partners (e.g. start-ups, business partners, outsourcing partners and research partners) had become increasingly important to achieve their strategic goals. Furthermore, transformational leadership (18 codings), and open communication and participation (18 codings) play key roles in empowering employees and being transparent in communication. Moreover, executives from 12 banks regard the setup of lean communication structures (12 codings) and the usage of social and digital media (12 codings), such as enterprise social networking, Twitter, Facebook, Webex, and so on, in their communication initiatives as a way to engage people. A few bank executives also mentioned mechanisms such as regular trainings and teamwork (9 codings), cross-functional trainings (7 codings), specific innovation rooms/meetings, and management dialogues and campaigns (6 codings) as improving the communication and transparency within their companies.

4.3 Control Variables

Finally, the following key patterns were elicited in relation to the observed control variables. ITG in German firms appears more agile than ITG in banks from Switzerland and Austria. Banks from Austria perceive their ITG as not very agile. Banks with an IT strategy focusing on digitalization for longer than four years view their ITG as very agile. Banks from Switzerland and Austria use both traditional dimensions (e.g. 15, 17, 37) and new dimensions (e.g. 7, 25, 26, 42). German and Austrian banks mainly engage in collaboration with external partners (dimension 47). Banks from Germany use new ITG processes (dimensions 43 and 47b). Very agile firms use new dimensions, including 10, 20, 44, and 47d. Large companies use new dimensions, including 7, 9, 20, 21, 23, 40, 44, and 47a. Large and very agile firms use the dimensions 20, 44 and 47. Small and very agile banks from Germany use dimension 1, 15, 30 and 47. This leads to the assumption, that the implementation of agile principles in ITG differs with respect to bank size and due to regional culture.

5 Discussion and Conclusion

The main objectives of this paper were to conduct a qualitative analysis to identify agile aspects of the governance construct and to elicit major patterns of ITG dimensions in the digital world. Uncovering the traditional and agile dimensions implemented in the banking industry to master digital innovations allows us to identify how the concept of an effective ITG framework has changed with regard to the demand for agility (RQ). As indicated in the “Data analysis” section, the analysis identified both traditional and agile ITG dimensions. Furthermore, it outlined agile characteristics of the uncovered agile mechanisms and offered several key ITG patterns.

Weil and Ross [11] illustrated the importance of implementing an effective ITG framework by means of the three major dimensions - structural, processual, relational - of ITG in an organization. They showed that up to 40% higher returns could be generated

by using effective ITG processes. However, because of sweeping changes in the new and fast-moving economy, leading to new requirements in and expectations of IT, the task of aligning business and IT remains a prime challenge, even though ITG frameworks are being implemented. Today's firms need flexible, complementary, adaptive and collaborative ITG dimensions to be put in place, if they are to prosper in a turbulent environment, in which the challenge for firms is to sustain value realization from IT instead of restraining its importance by emphasizing control. Therefore, what is required in the Digital Age is the underlying capability to flexibly adapt to changing business environments and requirements, to select and adopt promising new technologies, to effectively anticipate future needs, and to ensure that the potential residing in IT is proactively communicated within the firm, thus ensuring its effective exploitation [54]. Hence, the digital era demands a culture of speed and collaboration, if it is to differentiate and deliver extraordinary customer service to drive business growth. In this era, the rise of mobility and the velocity to deliver differentiated business processes is critical to success, which calls for a more agile ITG dimensions [14].

With regard to the understanding of the agile strategies within the governance construct, the respondents all believed that implementing agile strategies helps improve their agility. Therefore, 25 agile dimensions were identified. The most frequently employed agile mechanisms are the ITG dimension of the processes and the relational/communication mechanisms. Moreover, our study also indicates that traditional ITG dimensions are important for sustaining control. In this way, 22 traditional ITG dimensions were uncovered. In the context of the punctuated equilibrium theory, the concept of an effective ITG framework changing in response to the demand for agility in organizations calls for more ambidextrous approaches.

The two systems—traditional and agile—should work together, with a constant flow of information and activity between them [2]. In other words, to be effective an agile ITG governance must work seamlessly and organically with traditional, as well as agile ITG dimensions, so that the whole organization is both ensuring that tasks are completed with efficiency and reliability, constantly and incrementally improving itself, and handling today's increasingly strategic challenges with speed and agility. Therefore, the interaction between the traditional and the agile dimensions needs to be optimized in managing strategies to impact positively on the agility of a company.

Finally, this study is not without limitations. First, the scope of the data collected is restricted to firms from Germany, Switzerland and Austria. Considering other geographic locations could provide additional insights. Second, the analysis is limited to the banking industry. As such, understanding might be advanced from investigating other sectors. Third, we mainly interviewed bank executives. Use of different group, such as employees working in IT and business, could lead to further relevant insights.

In conclusion, our interview data provide relevant perspectives of traditional and new ITG dimensions implemented within the banking industry of the German-speaking part of Europe, which highlights the ambidextrous aspect within the governance construct. As such, the current study helps stimulate further investigation into combining agile capabilities with governance capabilities.

References

1. Collin, J., Halén, M., Helenius, M., Hiekkänen, K., Itälä, T., Korhonen, J.J.: IT Leadership in Finnish Organizations and Digital Transformation. Alto University (2014)
2. Kotter, J.P.: Accelerate Building Strategic Agility for a Faster-Moving World. Harvard Business Review Press, Boston (2014)
3. Peterson, R.R.: Integration strategies and tactics for information technology governance. In: van Grembergen, W. (ed.) *Strategies for Information Technology Governance*, pp. 37–80. IGI Global, Hershey (2004)
4. Luna, A.J.H., Costa, C.P., Moura, H.P.D., Novaes, M.A., do Nascimento, C.A.: Agile governance in information and communication technologies, shifting paradigms. *J. Inf. Syst. Technol. Manag.* **7**(2), 311–334 (2010)
5. Cheng, T.-H., Jansen, S., Remmers, M.: Controlling and monitoring agile software development in three Dutch product software companies. In: *Proceedings of the ICSE Workshop on Software Development Governance*, pp. 29–35. IEEE, Washington, D.C. (2009)
6. Qumer, A.: Defining an integrated agile governance for large agile software development environments. In: *Agile Processes in Software Engineering and Extreme Programming*, pp. 157–160 (2007)
7. Fruhling, A., McDonald, P., Dunbar, C.: A case study: introducing eXtreme programming in a US Government system development project. In: *Proceedings of the 41st Annual Hawaii International Conference on System Sciences (HICSS)*, Waikoloa, HI, USA, p. 464. IEEE (2008)
8. Almeida, N.H.R., de Magalhaes, E.M.C., de Moura, H.P., de Teixeira, F.J.G.D.A., Cappelli, C., Martins, L.M.F.: Evaluation of a maturity model for agile governance in ICT using focus group. In: *Proceedings of the Annual Conference on Brazilian Symposium on Information Systems*, 15–22. Brazilian Computer Society, Porto Alegre, Brazil (2015)
9. Schadewitz, N., Jachna, T.: Comparing inductive and deductive methodologies for design patterns identification and articulation. In: *International Design Research Conference IADSR 2007 Emerging Trends in Design Research*, Hong Kong (2007)
10. Bhattacharjya, J., Chang, V.: Adoption and implementation of IT governance. cases from Australian higher education. In: *Strategic Information Systems: Concepts, Methodologies, Tools, and Applications*, pp. 1308–1326. IGI Global, Hershey (2010)
11. Weill, P., Ross, J.W.: IT governance. How top performers manage IT decision rights for superior results. Harvard Business School Press, Boston, Mass (2004)
12. van Grembergen, W., de Haes, S.: Measuring and improving IT governance through the balanced scorecard. *Inf. Syst. Control J.* **2**(1), 35–42 (2005)
13. IT Governance Institute: Cobit 4.1: Framework, Control Objectives, Management Guidelines, Maturity Models. Meadows (2007)
14. Vejseli, S., Rossmann, A.: The impact of IT Governance on firm performance: a literature review. In: *Proceedings of the 21st Pacific Asia Conference on Information Systems (PACIS)*. AIS Electronic Library (2017)
15. de Haes, S., van Grembergen, W.: IT governance structures, processes and relational mechanisms: achieving IT/Business alignment in a major Belgian financial group. In: *Proceedings of the 38th Annual Hawaii International Conference on System Sciences (HICSS)*, p. 237b. Los Alamitos, CA (2005)
16. Sambamurthy, V., Zmud, R.W.: Arrangements for information technology governance. a theory of multiple contingencies. *MIS Q.* **23**(2), 261–290 (1999)

17. Weill, P., Ross, J.: A matrixed approach to designing IT governance. *MIT Sloan Manag. Rev.* **46**(2), 26 (2005)
18. de Haes, S., van Grembergen, W.: IT governance and its mechanisms. *Inf. Syst. Control J.* **1**, 27–33 (2004)
19. van Grembergen, W. (ed.): *Strategies for Information Technology Governance*. IGI Global, Hershey (2004)
20. Weill, P.: Don't just lead, govern. How top-performing firms govern IT. *MIS Q. Executive* **3**(1), 1–17 (2004)
21. Symons, C.: *IT governance framework*. Forrester Research (2005)
22. Burtscher, C., Manwani, S., Remenyi, D.: Towards a conceptual map of IT governance: a review of current academic and practitioner thinking. In: *Proceedings of the UK Academy for Information Systems Conference*, p. 15. AISeL (2009)
23. Almeida, R., Pereira, R., Mira da Silva, M.: IT governance mechanisms: a literature review. In: Falcão e Cunha, J., Snene, M., Nóvoa, H. (eds.) *IESS 2013. LNBIP*, vol. 143, pp. 186–199. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-36356-6_14
24. de Haes, S., van Grembergen, W.: An exploratory study into IT governance implementations and its impact on business/IT alignment. *Information Systems Management* (2009)
25. Gottschalk, P.: *E-business Strategy, Sourcing and Governance*. IGI Global, Hershey (2005)
26. Gersick, C.J.G.: Revolutionary change theories, a multilevel exploration of the punctuated equilibrium paradigm. *Acad. Manag. Rev.* **16**, 10–36 (1991)
27. Romanelli, E., Tushman, M.L.: Organizational transformation as punctuated equilibrium: an empirical test. *Acad. Manag. J.* **37**(5), 1141–1166 (1994)
28. Greiner, L.E.: Evolution and revolution as organizations grow. a company's past has clues for management that are critical to future success. *Family Bus. Rev.* **10**(4), 397–409 (1997)
29. Tushman, M.L., O'Reilly, C.A.: Ambidextrous organizations: managing evolutionary and revolutionary change. *California Manag. Rev.* **38**(4), 8–30 (1996)
30. Dunphy, D.C., Stace, D.A.: Transformational and coercive strategies for planned organizational change: beyond the O.D. model. *Organ. Stud.* **9**(3), 317–334 (1988)
31. Higgs, M., Rowland, D.: All changes great and small: exploring approaches to change and its leadership. *J. Change Manag.* **5**(2), 121–151 (2005)
32. O'Reilly, C.A., Tushman, M.L.: Organizational ambidexterity. past, present, and future. *Acad. Manag. Perspect.* **27**(4), 324–338 (2013)
33. Markides, C.C.: Business model innovation: what can the ambidexterity literature teach US? *Acad. Manag. Perspect.* **27**(4), 313–323 (2013)
34. March, J.G.: Exploration and exploitation in organizational learning. *Organ. Sci.* **2**(1), 140–145 (1991)
35. Markides, C., Charitou, C.D.: Competing with dual business models: a contingency approach. *Acad. Manag. Executive* **18**(3), 22–36 (2004)
36. Kotter, J.: How the most innovative companies capitalize on today's rapid-fire strategic challenges-and still make their numbers. *Harvard Bus. Rev.* **90**(11), 43–58 (2012)
37. van Roosmalen, M.W., Hoppenbrouwers, S.: Supporting corporate governance with enterprise architecture and business rule management: a synthesis of stability and agility. In: *Proceedings of the International Workshop on Regulations Modelling and Deployment (ReMoD 2008)*, pp. 13–24. CEUR-WS.org (2008)
38. Sloane, E., Beck, R., Metzger, S.: AGSOA - Agile Governance for Service Oriented Architecture (SOA) systems: a methodology to deliver 21st century military net-centric systems of systems. In: *Proceedings of the 2nd Annual IEEE Systems Conference*, pp. 1–4. IEEE (2008)

39. Luftman, J.N., Lewis, P.R., Oldach, S.H.: Transforming the enterprise: the alignment of business and information technology strategies. *IBM Syst. J.* **32**(1), 4–16 (1993)
40. Sambamurthy, V., Bharadwaj, A., Grover, V.: Shaping agility through digital options: reconceptualizing the role of information technology in contemporary firms. *MIS Q.* **27**(2), 237–263 (2003)
41. Sherehiy, B., Karwowski, W., Layer, J.K.: A review of enterprise agility: concepts, frameworks, and attributes. *Int. J. Indus. Ergon.* **37**(5), 445–460 (2007)
42. Yang, C., Liu, H.M.: Boosting firm performance via enterprise agility and network structure. *Manag. Decis.* **50**(6), 1022–1044 (2012)
43. Goldman, S.L., Nagel, R.N., Preiss, K.: *Agile Competitors and Virtual Organizations: Strategies for Enriching the Customer*. Van Nostrand Reinhold, New York (1995)
44. Luna, A.J.H.D.O., Kruchten, P., Moura, H.P.D.: GAME. Governance for Agile Management of Enterprises: a management model for agile governance. In: Proceedings of the 8th International Conference on Global Software Engineering workshops (ICGSEW), pp. 88–90. IEEE, Piscataway (2013)
45. Luna, A.J.D.O., Kruchten, P., Pedrosa, M.L.D.E., Neto, H.R., de Moura, H.P.: State of the art of agile governance: a systematic review. In: *IJCSIT* (2014)
46. Beck, K., Beedle, M., van Bennekum, A., Cockburn, A., Cunningham, W., Fowler, M., Grenning, J., Highsmith, J., Hunt, A., Jeffries, R., Kern, J.: *The agile manifesto* (2001). <http://agilemanifesto.org/>. Accessed 15 Nov 2017
47. Madi, T., Dahalin, Z., Baharom, F.: Content analysis on agile values: a perception from software practitioners. In: Proceedings of the 5th Malaysian Conference in Software Engineering (MySEC), pp. 423–428. IEEE, Piscataway (2011)
48. Fernandes, J.M., Almeida, M.: Classification and comparison of agile methods. In: Abreu, F.B.E. (ed.) Proceedings of the 17th International Conference on the Quality of Information and Communications Technology (QUATIC), pp. 391–396. IEEE, Piscataway (2010)
49. Benbasat, I., Goldstein, D.K., Mead, M.: The case research strategy in studies of information systems. *MIS Q.* **11**, 369–386 (1987)
50. Saunders, M., Lewis, P., Thornhill, A.: *Research Methods for Business Students: Financial Times*, 5th edn. Prentice Hall, Harlow (2009)
51. Saetang, S., Haider, A.: The impacts of IT governance implementation: a case study on banking industry in Thailand. In: Proceedings of the Portland International Center for Management of Engineering and Technology (PICMET), pp. 2619–2627. IEEE (2013)
52. Kuckartz, U., Dresing, T., Rädiker, S., Stefer, C.: *Qualitative evaluation. Der Einstieg in die Praxis*, 2nd edn. VS Verlag für Sozialwissenschaften/ GWV Fachverlage GmbH Wiesbaden, Wiesbaden (2008)
53. de Haes, S., van Grembergen, W.: *Enterprise Governance of Information Technology, Achieving Alignment and Value, Featuring COBIT 5, Management for Professionals*, 2nd edn. Springer, Cham (2015). <https://doi.org/10.1007/978-3-319-14547-1>
54. Schlosser, F.: Mastering the social IT/business alignment challenge. In: Proceedings of the 18th Americas Conference on Information Systems 2012 (AMCIS), pp. 1843–1849. AIS Electronic Library (2012)



Organizations in Transformation: Agility as Consequence or Prerequisite of Digitization?

Dominic Lindner¹(✉) and Christian Leyh²

¹ Chair of IT Management, Friedrich Alexander Universität Erlangen-Nürnberg,
Lange Gasse 20, 90403 Nuremberg, Germany

Dominic.Lindner@fau.de

² Chair of Information Systems, ESP IS in Manufacturing and Commerce,
Technische Universität Dresden, Helmholtzstr. 10, 01069 Dresden, Germany

Abstract. With the aim of meeting new technological and social requirements, nearly all companies are facing the so-called digital transformation. A starting point to respond to the changing demands is to improve the company's ability to operate flexibly and, thus, increase its agility. Thereby, with this paper, we contribute to the ongoing discussion on changing corporate organizational structures. We used a three-stage study design to provide insights on the question of whether agility is a prerequisite or a consequence of the digital transformation. In addition to a literature analysis, we present the results of focus group interviews with managers of selected enterprises.

Keywords: Digital transformation · Digitalization · Digitization · Agility
Corporate organizational structures

1 Companies in the Area of Tension of Digital Transformation

Nowadays enterprises are facing a complex environment with rapid and profound changes. According to Bajer [1], one of the main drivers of these changes is the exponential technological development, which is often strongly related to buzzwords like “digital transformation,” “digital age,” or “digital change.” These advanced technological opportunities result in new fundamental paradigm shifts and affect almost every enterprise, regardless of size or industry sector [2, 3]. This also applies to the German corporate landscape, which is predominantly characterized by small and medium-sized enterprises (SMEs) [4]. However, SMEs are often still classically organized and optimized for controllability and control, which is in contradiction to the fundamental variations and high velocity of digital changes [5]. To address this contradiction, our article focusses on SMEs (according to the SME definition of the European Union [6]). Thereby, traditional SMEs are usually characterized by shorter decision-making processes and greater flexibility for customers and employees, as well as smaller budgets than large companies [7]. According to Urbach and Ahlemann [5], the majority of these traditional SMEs are also strongly hierarchically structured and have clear communication structures oriented towards hierarchical relationships. The traditional SMEs are, therefore, optimized for control, monitoring, and coordination. In the course of digital

transformation, however, it is recommended that companies should be organized in a more networked and project-oriented manner [8–10]. To achieve this, it is necessary for a traditional SME to restructure its organization.

Here, the discussion about the organizational structures of companies has long been conducted in the literature. Thus, “Taylorism,” with its optimization of controllability and control, was initially the subject of discussion. At the turn of the millennium, in the course of the “humanization of work,” the discussion was extended to include agility in companies. Along with digitization, this has also been extended by a so-called digital business organization, which is controversially discussed. However, consensus has been reached in the literature that increasing agility can be achieved with, or at least should be supported by, the help of the technology. This has also been a major issue in practice. Since current traditional forms of organization in companies do not always achieve the desired/necessary success in the context of digital change, companies are also increasingly fostering agile and/or digital restructuring [3, 9, 10].

Thereby, part of this discussion is the role of agility in the field of digital change. This is shown by two contradictory viewpoints: Experts [11–13] argue that only increased agility in the corporate organizational structures enables and enforces the digitization and digitalization of a company, whereas scholars [3, 14, 15] argue that the implementation of new digital “components” is a prerequisite for increased agility in the organization. In any case, it becomes clear that agility in corporate organizational structures is a major aspect within the digital transformation of enterprises. With regard to SMEs, more specific conditions apply; SMEs are often less digitized due to small budgets but are often more agile because of their short decision-making processes. Therefore, a separate assessment of those two viewpoints for SMEs is necessary within this discussion. Accordingly, the aim of our paper is to examine the following research questions (RQs):

- RQ 1: *Is agility a prerequisite for the digital transformation within companies, especially within SMEs, or is it a result of digital transformation?*
- RQ 2: *What recommendations for action can be derived regarding agility and digital transformation, and what does an implementation approach look like?*

To address these research questions, we set up a three-stage study (see Fig. 1). This study consists of a systematic literature analysis, (focus) group discussions, and reference modeling. To present our study results, the paper is structured as follows. Following this introduction, the second section describes the overall design of our study in detail for each step. The third section deals with the results of the literature analysis and provides state-of-the-art knowledge concerning the discussion of agility in the organizational structures of enterprises. Sections four and five discuss the results of the focus group interviews, and section six describes the reference model for the implementation of agility in SMEs. The paper concludes with a brief discussion and summary, as well as an outlook for future research steps in our long-term research.

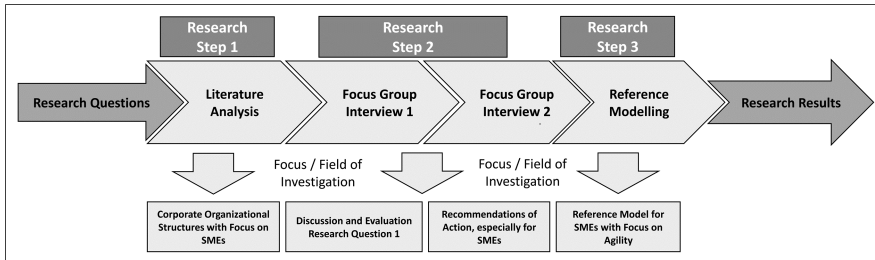


Fig. 1. Research Steps

Therefore, one aim of our article is to contribute to the discussion on enterprise organization changes in the context of digital transformation and agility with a specific focus on SMEs. More specifically, we discuss the interdependencies between agility and digitization from the practical view of SMEs. By addressing RQ 1 we provide an overview of existing literature about digitization and agility. Pilot projects described in the group discussions give a first insight in agility of SMEs and can be used as foundation for further research. Thereby, RQ 2 has a clear focus on the practitioners' viewpoint and by addressing RQ 2 we set up a framework for an agile organization.

2 Study Design

2.1 Research Step 1 – Literature Analysis

The aim of the literature analysis as the first step of our study (see Fig. 1) was to identify the state-of-the-art knowledge regarding RQ1. We followed the approach of Fettke [16]. These set of guidelines for a literature analysis includes literature searches using search strings in (academic) databases, literature evaluation, and a deeper analysis, followed by an interpretation of the review results [16]. Our literature evaluation step was also combined with methodology from Mayring [17] to classify the existing knowledge into specific categories (structuring content analysis).

For the literature analysis, we selected the databases “SpringerLink,” “Emerald,” “Business Source Complete,” and “ScienceDirect.” For the conducted search within the databases, we used the respective search functionality with a combination of the following search strings: (“company” OR “organization”) AND (taylor* OR agile* OR digi*) AND (“transformation” OR “framework”). Thereby, the search was limited to abstracts, titles and keywords.

Our initial search step identified 1,169 results. Although these papers were distributed between 1911 and 2017, the articles up to the year 2000 were only examined on a random sample basis in order to briefly illustrate the prior discussions on corporate organizational structures. In order to identify relevant articles between 2000 and 2017, the abstracts and the titles of the respective papers were examined for relevance to the research questions and, if necessary, the articles were read in-depth. This resulted in 50 contributions, which were examined in more detail. All 50 papers were then read in full and were evaluated in terms of our research focus. Thereby, 12 articles remained that

could be categorized (in accordance with the approach of [17]) in two categories: Six papers were categorized under “Agility,” and six papers were categorized under “Digital Organization.” In addition, to those two categories, we analyzed the papers for their relevance regarding SMEs in combination with our research focus. Ultimately, we identified six additional relevant papers. The list of all 18 papers is not included in this paper but can be requested from the first author.

2.2 Research Step 2 – Focus Group Interviews

The aim of the focus group interviews as the second step in our study was to evaluate and clarify the identified issues from the literature analysis and to derive recommendations for action for SMEs. The advantage of a focus group discussion lies in the collective decision making as well as the fact that, in contrast to individual interviews, new aspects of the research questions are discovered and discussed through the interaction of the group members [18, 19]. Altogether, we conducted two focus group discussions. In the first one (July 2016), the discussion was evaluated regarding whether agility is a prerequisite or a consequence of digitization, and in the second one (September 2016), recommendations for actions to increase agility were derived.

The group discussions were conducted according to the approach of Krueger and Casey [18]. The participants were selected based on existing contacts. In addition, we searched different “I offer” tags on the German business portal XING (www.xing.com). For the final selection, the potential participants had to be involved in an agile transition, had to hold a managerial position, and the company in which they were working had to be founded at least five years ago to ensure that the participants would be able to speak of a change of established structures. A table of all participants is attached in the Appendix (see Table 2). As shown, we involved members of the board of directors and managing directors of SMEs as well as large-scale companies and a member of the board of directors of a trade union in the group discussions. The participants of the large-scale companies were invited to bring further viewpoints into the discussion since those companies often use more and more advanced and therefore, it was possible for the participants of the SMEs to gain insight in those more “digitized” companies.

Both group discussions were recorded and transcribed. From this transcript, statements on the research questions were derived. For the first group discussion, participants were asked to present pilot projects in which they implemented “digitization aspects.” Four participants agreed to introduce their projects. Each pilot project was presented for 20 min, followed by 40 min in an open dialogue. In the second group discussion, recommendations for increasing agility in SMEs were determined based on the first group discussion. For both group discussions, a partial structured form was chosen. In accordance with recommendations from Greenbaum [19], the questions were formulated in advance and presented in a pre-established order; this procedure ensured that all questions were posed and adequately discussed and that the research questions and important components were focused.

2.3 Research Step 3 – Reference Modeling

The participants' recommendations for action were combined into a reference model in order to organize their ideas in a clear and concise manner for possible further research and use in practice. The reference model is motivated by RQ 2 to provide practitioners a clear visualization of the given recommendations of the group discussions.

Reference modelling is used to derive a simplified representation of reality from observations or theories [20]. A reference model merely provides a prefabricated solution scheme for a possible situation to deal with a practical problem, and does not necessarily have to be unrestrictedly valid in practice [21]. For the conception of a reference model, the practical problem is, first, recorded. Subsequently, a reference framework is created and the problem solution is presented. In our study, the research questions were set in a specific context described by a fictitious SME as a reference framework to discuss an explicitly practical problem. This SME is working for several clients and provides numerous services for them. It is organized by several departments and disciplines, like HR, Marketing, and Finance. Based on this fictitious SME, the participants of the second group discussion had to provide solutions and recommendations for action that were used for designing the reference model in the third step of our study.

3 Results of the Literature Analysis

This section shows the results of the literature analysis, which is aimed at highlighting the ongoing discussion on corporate organizational structures. Nevertheless, although this discussion expounds upon the two directions of agility and digitization, among other things, the connection between these two aspects has hardly been investigated, and that there is hardly any literature with reference to classic SMEs that address agility and digitization.

Since we are using in the following discussion the two central terms “agility” and “digitization” extensively, a short definition shall be given beforehand:

“Agility” means the ability of a company to be prepared for and to respond to changing capacity demands and changing functional requirements very quickly (if possible in real time), as well as the possibilities of the used information technology to extend the capacity if necessary. Thereby, agility and flexibility are not clearly differentiated [22]. Therefore, the assumption for our study is that agility leads to more flexibility.

The term “digitization” is defined according to Petry [13] as the (exponential) implementation and usage of (information) technology in companies. Thereby, digitization or digital transformation (as synonym) affects many sections and components of enterprises, even across company borders. Linking departments, as well as the contact to customers or suppliers, is important in the context of digitization. ICT triggers and enables this “transformation” of the company towards a holistic network. This means digitization is not only defined by a restructuring or optimization of the enterprises' ICT landscapes, though this is the main aspect of digitization [2].

3.1 Corporate Organizational Structures vs. Digitization

The discussion about the digitization of organizations is by no means new and does not directly concern (agile) organization structures. As early as 1985, Bair [23] investigated the change of organizations due to the personal computer (PC) and called for a more intense discussion on the digitization of organizations, which, however, has not shown any further noteworthy developments since the turn of the millennium. Nevertheless, this discussion has increased drastically in number of publications since 2012. The goal of the discussed digitization is often uniform in the identified literature: Urbach & Ahlemann [5] as well as Schrauzer [15] believe that digitization is the precursor to increased agility in organizations and their structures. In their studies, Bley et al. [3] showed that the digitization of SMEs is increasing significantly and that these new possibilities open more agile ways of responding to customer requirements. Martensen et al. [14] focused on the computerization of workplaces and noted that a more flexible reorganization of companies is possible through mobile hardware and software. They concluded that the digitization is an enabler for new and more flexible forms of business. Noonan et al. [24] concretized this and showed in their framework that the increasing flexibility of organizations is a prerequisite for the necessary reorganization of internal IT in the current digital age.

Thus, the majority of identified sources support the view that increasing digitization of companies can lead to more agile corporate organizational structures.

3.2 Corporate Organizational Structures vs. Agility

In contrast to the papers from Sect. 3.1, there is also a discussion with more focus on the agility of companies. Gastaldi et al. [25] believe that the required IT reorganization in the digital age can demand disruptive changes from companies. They stated that new digital technologies cannot be implemented in a classical/Taylorist organization and demanded more agile structures. Arnold [26] stated a similar opinion and described the example of “Haufe-umantis,” which use different methods like the democratic vote of the leaders and self-organized teams. To attain a higher agility, Arnold [26] and Jovanović et al. [27] recommended experiments to determine how higher agility can be achieved in combination with the digital transformation.

Lindner et al. [10] recommended adapting organizations to the vision of the New Work (see [28]): Work when, how, and where you want. In this way, employees are given maximum freedom in the execution of their work duties. Lindner et al. [10] noted that the flexibilization of work is strongly dependent on mobile technology, but that the necessary organizational restructuring in terms of more agility should be undertaken before new technology is implemented as this facilitates the implementation of technology. Furthermore, the potential that new technology offers cannot be achieved sufficiently without appropriate agile organizational structures [11, 12].

The majority of the sources identified in this section, thus, support the viewpoint that increasing agility in the company is a precursor to successful digitization.

3.3 Agility and Digitization in SMEs

The papers identified were also examined in the context of an SME focus. Similar to the discussion on digitization in Sect. 3.1, case studies on the digitization of an SME, such as [8], can be found showing an increase in flexibility and networking through the additional usage of new software tools, such as data analysis tools (see also [29]). Within the SME focus, the smaller budget and the fact that SMEs are often less digitized than large companies have been addressed. For this reason, a centralized and rather inflexible IT system landscape is often found in SMEs, which, due to the smaller budget, is often built up step by step and over a long period of time [30].

In contrast, in the papers focusing on agility, SMEs, compared to large companies, are often seen and described as more agile and flexible due to their shorter decision-making processes [8]. However, many SMEs are subject to constant and often rapid growth, which can reduce their flexibility and agility. Therefore, Branicki et al. [9] were investigating how the so-called “entrepreneurial spirit” of SMEs can be maintained and how agility can be better scaled in an SME. Ng and Kee [31] also believe that maintaining or increasing agility is an important factor for the growth of an SME.

In summary, SMEs are considered to have lower levels of digitization, which is to be gradually increased and often have less modern centralized IT. Nevertheless, SMEs are described as very flexible and agile. It is, therefore, important to maintain this agility even when facing rapid growth.

3.4 Discussion and Summary of the Literature Analysis Results

The discussion addressing corporate organizational structures has been long debated. Some of the authors believe that digitization with flexible hardware and software can increase agility. For them, it is also not possible to work in a more flexible or agile way in terms of time and place without appropriate technology. However, other authors believe the implementation of technology requires increased agility beforehand, and that the required reorganization of IT in the digital age is only barely possible in enterprises with traditional organizational structures. This, again, shows the topicality of the question of whether agility is a prerequisite or a consequence of digitization.

A similar discussion is taking place with regard to SMEs, but with a slightly different context. For example, SMEs often have a more centralized and less developed IT system landscape than large companies due to their lower budget, and it is often necessary to expand this step by step. SMEs are already described as very agile in contrast to large corporations; therefore, it is necessary to maintain or increase their agility, especially when SMEs face rapid growth. However, even if the question of agility vs. digitization has not been clearly answered yet in the literature, the SME focus still has not been addressed in the literature to a large extent. Hence, due to the specific characteristics of the SME and their importance for the economy, it is essential to discuss these aspects with a special focus on SMEs. Therefore, in the following sections, we show and discuss the results of our focus group interviews.

4 Results of Group Discussion 1

The aim in the first group discussion was to examine current agile or digital pilot projects and to consider, in a mutual dialogue, whether agility is a prerequisite or consequence of digitization. In our study, a “pilot project” is a company project where a company implements and tries out new technology or working method with one single team. The goal of a pilot project is to evaluate if the new technology/method increases the efficiency of the team and leads to a benefit for the company.

4.1 Agility as a Precondition for Digitization

In order to evaluate the theses, four of the participants (No. 1, 2, 3, and 4, Table 2) presented their current pilot projects. This assisted in enabling the deduction of whether they see agility as a prerequisite for an increased agility or as a consequence. In this section, the examples are categorized. They address agility on the basis of employee satisfaction as well as agility on the basis of the complexity of digitization.

Agility with Respect to the Employees

The first contribution from the group discussion results from the current focus of one of the largest German trade unions. Many of the union’s member companies are currently aiming for a digital transformation. According to the union’s executive board member (No. 1, Table 2): “Digitization matters everywhere. [...] And in this respect, more and more agile methods are inevitably being worked on.” Consequently, the executive board addressed a necessity to use agile methods. The executive board member also claimed that “agile methods are good for the personnel,” which was especially justified with the increased flexibility in work execution.

Examples of such pilot projects, according to the union’s executive board member, are the establishment of an agile team as a so-called “u-boot project” or internal project, equipping teams with new technologies not available yet in the company as standard, as well as the acquiescence of home offices and flexible work locations.

Another contribution comes from an IT service provider of CRM software (No. 2, Table 2) and is based on the motivation of the company to become attractive to skilled employees. Numerous new digitization projects on the customer side require new specialists. This pilot project is based on the beliefs of the managing director, that agility can increase employee satisfaction. “[We have a] location disadvantage. [Our location] is a beautiful city, but seated at the suburbs of Stuttgart. That means all good employees are attracted to [the large concerns of this region],” the managing director said. In order to test various ways to increase agility, the members of the management board restructured the company and reorganized the business units as follows:

- Holacracy (see [32]) for the “Consulting” division
- Scrum (see [33]) for the “In-house software” division
- Classical hierarchy for the “Software solutions” division

The managing director emphasized that he wanted to offer each employee an appropriate method and “[create] a creative environment [with] regulations and framework

conditions, which allow us to have this flexibility.” The participants of the group discussion agreed that the demands of digitization projects for employees necessitate more flexibility in their work execution. According to the consensus of the participants, rigorous processes and limited flexibility encourage negative dynamics, such as dissatisfaction, burnout, and demotivation. They also agreed with the union’s executive board member, that current employment agreements are not suitable for the required flexibility requirements of agile work. To sum up, all participants agreed that the guidelines for an agile organization are an important prerequisite for digitization.

Agility with Respect to the Customer

The marketing manager (No. 3, Table 2), a member of the management board of an IT logistics service provider, presented the approach of a “more agile business relationship” with a marketing agency. In the approach the logistics experts in the company should give the cooperating marketing agency the vision of the task. Namely, in this case: “Create the entry into digital marketing with us. We wanted to do it that way consciously. We did not want to say: Make us a website. Make us an SEO campaign, make us somehow a Google AdWords thing. But step in with us.” The central question from the marketing manager was: “How do I describe the targets and tasks in such a dynamic environment? [...] Agile methods are good for that.” The consensus of the participants is that the change processes in the context of digitization ensure a high complexity and cannot be successfully realized without agile methods.

The participants summarized and thereby realized that, without increased agility, certain digitization projects would not have been possible in their companies. They agreed that an increased agility is important for the digitization of a company.

Another pilot project, where agility is also implemented because of numerous digitization projects, was presented by a member of the management board of an e-commerce service provider (No. 4, Table 2). Among other things, the e-commerce service provider is responsible for the digitization of customers’ sales processes. However, these projects often led to complaints and could no longer achieve the desired success. The manager considered a lack of agile principles to be the cause of many customer complaints and stated: “So far, it has often been [...] very waterfall-like. No feedback, no customer contact, no stakeholder contact.” The manager also noticed that many orders were made exactly according to the customer’s requirements, but this is not always optimal: “Many of the requirements were written as follows: ‘In menu X, I would like a button for Y.’ That is how the team got it and did it. [...] It did not work at all.” For the implementation, the manager decided to establish the principles of Scrum (see [33]) in the company. His aim was to increase employee and customer satisfaction, as well as productivity. Thereby, he underlined the opinion of the marketing manager (No. 3, Table 2), that agile methods are a better way to implement digitization projects. In an open dialogue, the manager (No. 4, Table 2) summarized: “Well, I can now say after six months: Creativity has come into being. The contact [with the customer] is always there.”

4.2 Discussion and Summary

All participants were trying to increase agility due to upcoming digitization projects. The consensus regarding the presented pilot projects was that increased agility is a prerequisite for more successful digitization. None of the participants who presented the projects saw digitization as a consequence of agility, which is noteworthy as many authors are pursuing this thesis. As shown in the literature review, the flexibility of place and time can be achieved especially with more mobile technologies.

To sum up, the four examples are divided into two argumentations concerning why increased agility before digitization seems important for companies. Therefore, the following theses justifying the participants' decisions can be derived:

- **Statement 1:** Digitization projects in companies impose increased flexibility demands on employees in the work execution. Agility can provide greater flexibility in the work execution as well as increased employee motivation and productivity. In addition, new employees will be needed, for which a more attractive working environment can be created by means of agile methods.
- **Statement 2:** The participants believe that digitization projects, in particular, bring an increased complexity, and that customer inquiries can be addressed more individually. Agile methods or principles are also seen as an efficient way to respond to digitization projects more flexibly and faster.

5 Results of Group Discussion 2

After the first group discussion showed that increased agility is a basic precondition for digitization, action recommendations for agility enhancement should be conceived to derive a reference model in the second group discussion. The question for the group discussion was, "How would you increase agility in a mature business like yours?" and the participants were then allowed to freely respond. The group discussions took place at intervals of almost three months.

5.1 Field of Tension for Scaling Agility

According to the participants, the central field of tension concerns how agility should be scaled in the company. The executive board member of the consulting company (No. 8, Table 2) stated: "Because we have been in the company for thirty years, [it means:] I have always done it this way. On the other hand, we have things that are extremely new, and finding the right balance is quite an exciting thing." Therefore, the participants recommended establishing different degrees of agility in the company. The executive board member of the data center (No. 5, Table 2) emphasized the following statement: "We have highly creative areas where cooperation must constantly reinvent itself, every day, as new requirements appear. But on the other hand, the classic data center, where it is really about that: one has to do some electrics that is repetitive." In this context, as a degree of agility the participants understand the need of clear processes and self-organization of employees. For example, high agility means less control, less processes

restrictions and high self-organization of employees while low agility means strict processes as well as a command and control leadership.

In addition, not every employee wants to work agilely, added the manager of the pharmaceutical company (No. 10, Table 2), who quoted an employee: “I do my job! Tell me what to do! And then I go home!” Thus, two key findings in the conception of the recommended actions were foregrounded:

- Not every employee wants to work agilely.
- Not every area of business needs agility.

5.2 Recommended Actions to Increase Agility

The recommendations for action are grounded on the two assumptions that not every employee wants to work agile and that not every department needs high agility. In the examples from the first group discussion, the participants highlighted that through increased agility, they seek flexibility in the work execution, increased employee satisfaction, reaction to change, and higher customer satisfaction. Therefore, the participants derived the following recommendations (see Table 1) for action from their pilot projects and experiences.

Table 1. Recommended actions from the focus group discussion

Requested attributes	Recommended actions within the reference model
Increased employee satisfaction	Areas of varying agility create possibilities for every employee in the company to choose freely
Flexibility in work execution	Greater autonomy in the work execution can give employees greater latitude to perform complex tasks
Faster reaction to change	By loosening the general framework, pilot teams can respond more effectively to changes in dynamic projects
Increased customer satisfaction	By direct customer contact, demands can be evaluated more quickly and change requests can be swiftly coordinated with customers

6 Reference Model to Enhance Agility

In this section, based on the recommendations for action, a reference model is formed to increase agility in a classic and fictitious SME. As indicated previously, since a reference model is not always fully applicable in practice nor can be presented in a simplified scheme, a limitation has been subsequently constructed.

The model in Fig. 2 is a summary of the participants’ recommendations for action in the context of agile SME transformation. It can be used as the basis for a first consideration to enhance agility. In order to display the demanded degree of agility, the participants recommended a division into three sectors differing in the proximity to the market. Outside the circle is the market. Thus, the central services often do mostly internal tasks and do not need to contact the customer or the market. Larger stable departments (e.g., addressing stable and long-term customers) with a mix of standard and individual

services require medium agility, and new pilot teams, in particular, require a high degree of agility. A pilot team (startup) can be separated from a department. Startups or pilot projects are projects that are very close to the market and often use agile methods – the closer to the market, the greater the autonomy in the work execution or the need for agility. With these different degrees of agility, the participants wanted to create a larger flexibility for employees. As such, agility should also be used purposefully. It is beyond question that agility may even be counterproductive in accounting. This can be justified due to strict regulations and standards. Outside are the departments of the company, which, as previously described, split from a certain size and form so-called startups. The departments are supported by the central services through standardized services such as HR, accounting, etc. As in the previous passage, in the absence of agility, these should be divided into smaller units. However, this is not easy, and it is currently not possible to specify exactly when a split or an integration should occur. The manager of the pharmaceutical company (No. 10, Table 2) stated: “You can do it well with the startups. But then it is hard to manage to recognize the point when it no longer works as a startup and you have to create other structures and mechanisms.” According to the participants, the startups do not necessarily have full-time employees, but rather employees can be assigned to a department and to one or more startups. The participants saw the basis for a company as a mix of classical and agile methods, depending on the area.

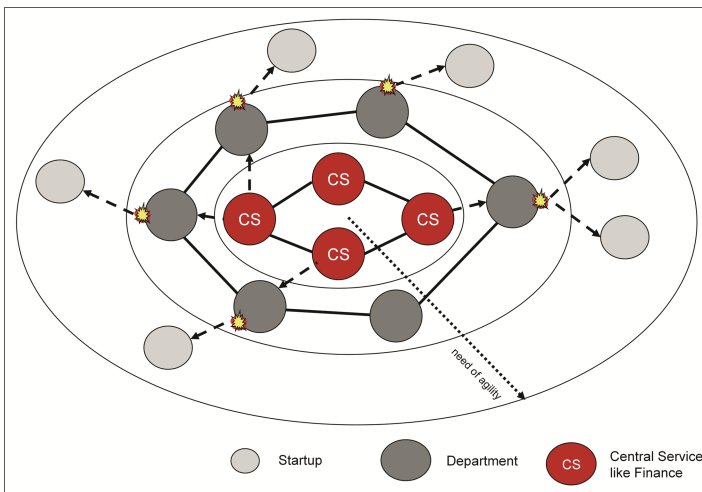


Fig. 2. Reference model based on the recommendations of the focus group participants

In summary, implementing these different methods in the company increases employee satisfaction. Furthermore, the autonomy in the work execution in startups fosters increased flexibility, which helps to better react to changes. In this way, customer satisfaction can be increased. This example considers a fictive SME, regardless of the industry or the market. However, the necessary change process behind the recommendations was not considered. Employees must be involved in the process, and legal aspects, such as employment contracts, etc., must be adapted. In addition, compliance

guidelines or technical requirements that might be necessary to connect many teams were not discussed, since this example focused on the creation of a first model. The control and the role of executives in the agile teams were also open questions, which could be examined in possible follow-up contributions.

7 Summary and Conclusion

The current discussion about business organization shows that, according to the Taylorist approach, a discussion about the humanization of work has occurred, which also covers the topic of agile organization. In the face of the exponential digitization, this debate is still tackling a general question: Is agility a prerequisite or consequence of digitization? Experts have not reached an agreement here.

In the first group discussion, all participants preferred the thesis stating that an increased agility is a prerequisite for digitization. Thus, the participants of the group discussion are currently trying to agilely transform their companies and increase their attractiveness for skilled workers, customers, and service providers through agile pilot projects. In the second group discussion, the participants developed concrete recommendations for action in the form of a model dividing the company systematically and in a self-organized manner, into small units. The classification is made according to the degree of required agility. Areas of high agility are subject to loosened framework conditions, greater autonomy in task completion and have direct customer contact. In this way, employee and customer satisfaction can be enhanced. Based on these recommendations, an abstract reference model could be designed and considered as a recommendation to agilely convert a company regarding the digital transformation.

As limitation, our study is based on a small sample of participants and as already stated, the reference model is only a simplified representation of reality and not always fully applicable. However, although these results represent only the views of a small group, it still broadens and continues the discussion on corporate organizational structures in the digital age. Thus, the question that remains unanswered is:

How, taking into account the smaller budgets of SMEs, can a meaningful organizational and IT architecture be designed to promisingly continue the increase of agility and flexibility of SMEs?

In our future research we will foster this unanswered question. Besides this question there are many other interesting fields of digitization of SMEs, e.g., the role of leadership in this kind of organizations. It would be interesting to understand how fast changing and self-organized teams can be controlled and motivated. Therefore, we recommend discussing current knowledge which is summarized under the term “digital leadership”. Also, we are focusing on the aspect how classical work-models will change with the increase of self-organization in companies (e.g., place and time flexible work). First approaches can be found with the terms “newwork,” “work 4.0” or “smart working.” In addition, we aim to further analyze our results by focusing on literature in strategic management that has provided plenty of insights regarding the appropriate balance of stability and flexibility of organizational structures, processes and work practices.

Appendix

Table 2. Participants of the focus group discussions

	Company	Participant's position	No. of employees	Company age in years	Participated in	
					Discussion 1	Discussion 2
1	Trade Union	Executive Board Member	120.000	70	X	
2	IT Service Provider (CRM)	Managing Director	120	15	X	
3	IT Service Provider (Logistics)	Managing Director	400	30	X	
4	IT Service Provider (E-Commerce)	Managing Director	30	5	X	X
5	IT Service Provider (Data Center)	Executive Board Member	200	25	X	X
6	IT Service Provider (Finance)	Senior Manager	7.000	50	X	X
7	IT Service Provider (Media Company)	Managing Director	1.600	30	X	
8	IT Service Provider (Consulting)	Executive Board Member	120	30	X	X
9	IT Service Provider (Industrial Sector)	Supervisory Board Member	200	10	X	X
10	Pharmaceutical company	Senior Manager	135.000	20		X

References

1. Bajer, J.: Digital transformation needs the human touch. *Strateg. HR Rev.* **16**, 91–92 (2017)
2. Mathrani, S., Mathrani, A., Viehland, D.: Using enterprise systems to realize digital business strategies. *J. Enterp. Inf. Manag.* **26**, 363–386 (2013)
3. Bley, K., Leyh, C., Schäffer, T.: Digitization of German enterprises in the production sector – do they know how “digitized” they are? In: *Proceedings of the 22nd Americas Conference on Information Systems, AMCIS 2016* (2016)

4. Ludwig, T., Kotthaus, C., Stein, M., Durt, H., Kurz, C., Wenz, J., Doublet, T., Becker, M., Pipek, V., Wulf, V.: Arbeiten im mittelstand 4.0 – KMU im spannungsfeld des digitalen wandels. *HMD Prax. Wirtsch.* **53**, 71–86 (2016)
5. Urbach, N., Ahlemann, F.: Der Wissensarbeitsplatz der Zukunft: Trends, Herausforderungen und Implikationen für das strategische IT-Management. *HMD Prax. Wirtsch.* **53**, 16–28 (2016)
6. European Union: Commission Recommendation of 6th May 2003 concerning the definition of micro, small and medium-sized enterprises. *Off. J. Eur. Union* 20th May 2003. (2003/361/EC), pp. 36–41 (2003)
7. Lindner, D., Ludwig, T., Amberg, M.: Arbeit 4.0 – Konzepte für eine neue Arbeitsgestaltung in KMU. *HMD Praxis Der Wirtschaftsinformatik* **6**(1), 17 (2018)
8. Arbussa, A., Bikfalvi, A., Marquès, P.: Strategic agility-driven business model renewal: the case of an SME. *Manag. Decis.* **55**, 271–293 (2017)
9. Branicki, L.J., Sullivan-Taylor, B., Livschitz, S.R.: How entrepreneurial resilience generates resilient SMEs. *Int. J. Entrep. Behav. Res.* (2017)
10. Lindner, D., Ott, M., Leyh, C.: Der digitale Arbeitsplatz – KMU zwischen Tradition und Wandel. *HMD Prax. Wirtsch.* **54**, 900–916 (2017)
11. Denning, S.: The next frontier for agile: strategic management. *Strategy Leadersh.* **45**, 12–18 (2017)
12. Beaumont, M., Thuriaux-Alemán, B., Prasad, P., Hatton, C.: Using agile approaches for breakthrough product innovation. *Strategy Leadersh.* **45**, 19–25 (2017)
13. Petry, T.: *Digital Leadership: Erfolgreiches Führen in Zeiten der Digital Economy*. Haufe Gruppe, Freiburg München Stuttgart (2016)
14. Martensen, M., Ryschka, S., Blesik, T., Bick, M.: Collaboration in the consulting industry: analyzing differences in the professional use of social software. *Bus. Process Manag. J.* **22**, 693–711 (2016)
15. Schrauzer, S.: Computerisierung in der globalen softwareentwicklung – Eine arbeitsmethodische betrachtung. *HMD Prax. Wirtsch.* **53**, 42–54 (2016)
16. Fettke, P.: State-of-the-art des state-of-the-art: eine untersuchung der forschungsmethode “Review” innerhalb der Wirtschaftsinformatik. *Wirtschaftsinformatik* **48**, (2006)
17. Mayring, P.: Qualitative Inhaltsanalyse. In: Mey, G., Mruck, K. (eds.) *Handbuch Qualitative Forschung in der Psychologie*, pp. 601–613. VS Verlag für Sozialwissenschaften, Wiesbaden (2010)
18. Krueger, R.A., Casey, M.A.: *Focus Groups: A Practical Guide for Applied Research*. Sage, Thousand Oaks (2015)
19. Greenbaum, T.L.: *The Handbook for Focus Group Research*. Sage, Thousand Oaks (1998)
20. Wilde, T., Hess, T.: *Methodenspektrum der Wirtschaftsinformatik: Überblick und Portfoliobildung*. Institut für Wirtschaftsinformatik und Neue Medien der Ludwig-Maximilians-Universität München (2006)
21. Schütte, R.: *Grundsätze ordnungsmäßiger Referenzmodellierung*. Gabler Verlag, Wiesbaden (1998)
22. Termer, F.: *Determinanten der IT-Agilität: Theoretische Konzeption, empirische Analyse und Implikationen*. Springer Fachmedien, Wiesbaden (2016)
23. Bair, J.H.: Personal computers and the office of the future. *Telemat. Inform.* **2**, 113–117 (1985)
24. Noonan, M., Richter, G., Durham, L., Pierce, E.: Learning and the digital workplace: what? so what? now what? *Strateg. HR Rev.* **16**, 267–273 (2017)
25. Gastaldi, L., Corso, M., Raguseo, E., Neirotti, P., Paolucci, E., Martini, A.: Smart working: rethinking work practices to leverage employees’ innovation potential. In: *Proceedings of the 15th International CINet Conference* (2014)

26. Arnold, H.: Digitalisierung der Unternehmensführung: Fallstudie und Plädoyer für mehr Mut in einer flachen Führung und Organisation. *Z. Führ. Organ. ZfO*. **85**, 330–336 (2016)
27. Jovanović, M., Mesquida, A.-L., Mas, A., Lalić, B.: Towards the development of a sequential framework for agile adoption. In: Mas, A., Mesquida, A., O'Connor, Rory V., Rout, T., Dorling, A. (eds.) *SPICE 2017. CCIS*, vol. 770, pp. 30–42. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-67383-7_3
28. Hackl, B., Wagner, M., Attmer, L., Baumann, D.: *New Work: Auf dem Weg zur neuen Arbeitswelt*. Springer Fachmedien, Wiesbaden (2017)
29. O'Connor, C., Kelly, S.: Facilitating knowledge management through filtered big data: SME competitiveness in an agri-food sector. *J. Knowl. Manag.* **21**, 156–179 (2017)
30. Essers, M.S., Vaneker, T.H.J.: Design of a decentralized modular architecture for flexible and extensible production systems. *Mechatronics* **34**, 160–169 (2016)
31. Ng, H.S., Kee, D.M.H.: Entrepreneurial SMEs surviving in the era of globalisation: critical success factors. In: Sindakis, S., Theodorou, P. (eds.) *Global Opportunities for Entrepreneurial Growth: Coopetition and Knowledge Dynamics within and across Firms*, pp. 75–90. Emerald Publishing Limited (2017)
32. Robertson, B.J.: *Holacracy: The New Management System for a Rapidly Changing World*. Macmillan Publishers, New York (2015)
33. Gloger, B.: *Scrum: Produkte zuverlässig und schnell entwickeln*. Hanser, München (2016)



Information Security Management Systems - A Maturity Model Based on ISO/IEC 27001

Diogo Proença^{1,2(✉)} and José Borbinha^{1,2}

¹ Instituto Superior Técnico, Universidade de Lisboa, Lisbon, Portugal
{diogo.proenca, jlb}@tecnico.ulisboa.pt

² INESC-ID - Instituto de Engenharia de Sistemas e
Computadores Investigação e Desenvolvimento, Lisbon, Portugal

Abstract. An Information Security Management System, according with the ISO/IEC 27001 is the set of “that part of the overall management system, based on a business risk approach, to establish, implement, operate, monitor, review, maintain and improve information security”. ISO/IEC 27001 defines the requirements and process for implementing an Information Security Management System. However, implementing this standard without a detailed plan can become a burden on organizations. This paper presents a maturity model for the planning, implementation, monitoring and improvement of an Information Security Management System based on ISO/IEC 27001. The purpose of this model is to provide an assessment tool for organizations to use in order to get their current Information Security Management System maturity level. The results can then be used to create an improvement plan which will guide organizations to reach their target maturity level. This maturity model allows organizations to assess their current state of affairs according to the best practices defined in ISO/IEC 27001. The maturity model proposed in this paper is evaluated through a multi-step perspective that is used to confirm that the maturity model makes a useful and novel contribution to the Information Security Management domain by taking in consideration the best practice of the domain.

Keywords: Information Security Management · Maturity model · Measurement Performance · Design

1 Introduction

In a growing and overly competitive world, only organizations that take advantage of the benefits the best information can deliver for decision-making are able to profit and thrive. Organizations should understand that information is such a valuable asset that it must be protected and managed properly. Information security should be used as a way to protect information against loss, exposure or destruction of its properties. [1] One of the goals of information security is to ensure business continuity while minimizing the impact of security incidents. In this sense, information is an asset that, like any other important asset, is essential to an organization and therefore needs to be adequately protected. This is especially important in the increasingly interconnected business

environment. As a result of this incredible increase in interconnectivity, information is now exposed to increasing numbers and a wide range of threats and vulnerabilities [2]. Information can exist in several forms. It can be printed or written on paper, electronically stored, transmitted by mail or by electronic means, presented in films or spoken in conversations. Whatever form is presented or the medium through which information is shared or stored, it is recommended that it be always adequately protected [2]. Information security is the protection of information from various types of threats to ensure business continuity, minimize risk to business, maximize return on investment and business opportunities. Information security is achieved by implementing a set of appropriate controls, including policies, processes, procedures, organizational structures, and software and hardware functions. These controls need to be established, deployed, monitored, critically reviewed and improved where necessary to ensure that the organization's business and security objectives are met. This should be done in conjunction with other business management processes [2]. Information security should always serve three elements [3]. The first is confidentiality, when we talk about confidentiality, we are talking about secrecy. Preserving the confidentiality of information means ensuring that only those who should have knowledge about it can access it. The second is integrity, the preservation of integrity involves protecting information against changes in its original state. These changes can be both intentional and accidental. The third and final one is availability, which ensures that information is accessible when someone who needs it tries to get it. The requested information must be provided as expected by the user.

The goal of this paper is to develop an artifact (a maturity model) by using an established approach to contribute to the Information Security Management body of knowledge. As a result, Design Science Research (DSR) was chosen as it combines the practical dimension and the scientific dimension. The maturity model focuses on the ISO/IEC 27001, which prescribes the requirements and process for implementing an Information Security Management System (ISMS), to define a maturity model for ISMS. In this paper we target our attention in answering two research questions (RQ), as follows:

- RQ1 - What are key requirements for an Information Security Management System process according to the ISO/IEC 27001 relevant for the purpose of maturity assessment?*
- RQ2 - How could a maturity Model specific to ISMS be designed which targets the challenges of different organizations and industries?*

To address these research questions, this paper is structured in six sections. First, the key terms and concepts are explained in Sect. 2. This is followed by Sect. 3, where the research methodology is outlined. Section 4, details the findings from a literature review in existing Information Security Management Maturity models and a comparison between the existing maturity models for the Information Security Management domain. Then Sect. 5, presents the ISMS Maturity model and the iterative development method used. The evaluation of the ISMS Maturity Model is presented in Sect. 6 which evaluates the mapping between the ISMS Maturity Model dimensions and the ISO/IEC 27001 requirements. This section also details the results of five assessments performed to five

different organization using the proposed maturity model. Finally, Sect. 7 details the conclusions and the limitations of the ISMS maturity model.

2 Foundation

This section explains the key terms and concepts within this paper, such as, “maturity models” and “information security management system” to ensure a common understanding.

In 1986, the US Department of Defense needed a method to assess the capabilities of the software companies with whom it worked, so Watts Humphrey, the SEI team and Miter Corporation were tasked with this task. In 1991 was released the first version, the CMM maturity model of capabilities. This model has achieved remarkable success and has been revised and improved having evolved into CMMI, the currently integrated capability maturity model integration version 1.3 [4].

Due to the success obtained, the principles used to develop the SEI maturity models served as inspiration to other authors, both academics and practitioners, and there are now hundreds of models applied to different domains [2]. Currently, the two major references of maturity models are CMMI and ISO/IEC 15504, both of which are related to Software Engineering processes.

In general, maturity can be defined as “an evolutionary progression in the demonstration of a specific skill or in the achievement of an objective from an initial state to a desired final state” [5]. In addition to the general definitions, there are many definitions of maturity that are directly related to the domain to which this term refers. As this work will develop a maturity model applied to a process of ISMS, it is also important to define maturity applied to a process. Maturity can then be defined as the “degree to which an organization executes processes that are explicitly and consistently documented, managed, measurable, controlled, and continuously improved. Maturity can be measured through appraisals” [4]. According to Loon [15], a maturity model is a sequence of maturity levels for certain objects, usually people, organizations or processes. In these models is represented the evolutionary path, anticipated, desired or typical, through discrete levels. In addition to the above, these models provide the necessary criteria to reach each of the model’s maturity levels. Thus, maturity models allow us to see at what level of the evolutionary process certain objects meet. The maturity levels are organized from an initial level of lower capacity to an advanced level corresponding to the maximum capacity of the reality in question. In order to reach higher maturity levels, it is necessary that there is a continuous progression of the capability of a given object.

ISO/IEC 27001 was based on the British standard BS7799 and ISO/IEC 17799. It was prepared to provide the requirements to establish, implement, operate, monitor, critically analyze, maintain and improve an ISMS [2]. An ISMS as defined by this standard is “that part of the overall management system, based on a business risk approach, to establish, implement, operate, monitor, review, maintain and improve information security” [2]. This standard is used around the world by all types of organizations as the basis for the organization’s policy management and implementation of information security. It is being used by small, medium and large organizations. In fact,

ISO/IEC 27001 is designed to be flexible enough to be used by any type of organization. This standard adopts the Plan-Do-Check-Act (PDCA) model, as depicted in Fig. 1, which is applied to structure all the ISMS processes.

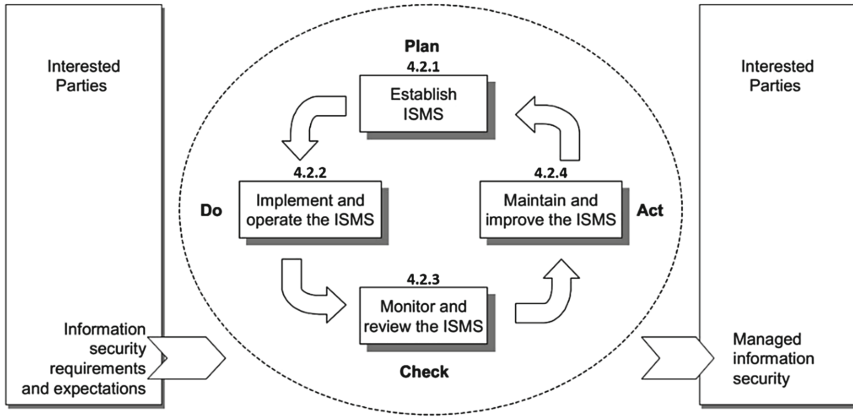


Fig. 1. PDCA model applied to ISMS processes and ISO/IEC 27001 mapping [2]

3 Research Methodology

In order to address the research questions of this paper, we selected the DSR paradigm [17, 19]. DSR is described by “a designer answering questions relevant to human problems via the creation of innovative artifacts, thereby contributing new knowledge to the body of scientific evidence. The designed artifacts are both useful and fundamental in understanding that problem” [19]. The major benefit of DSR is the fact that it addresses real-world problems and simultaneously contributes to the body of knowledge [17]. However, the development of maturity models within the Information Security Management domain is not new but has been popular for quite some time [6]. Mettler, et al. [12] count more than 100 models in the information systems domain, Poepplbuss et al. [14] counts even many more. One significant fault within this research area is the lack of specific contributions regarding how to develop maturity models. Moreover, most authors rarely describe their development process. Up to our knowledge there are only a few development procedure models for maturity models. The models of Becker et al. [16] and De Burin et al. [13] seem to be quite popular among the community based their citation counts. We decided to apply the model of Becker et al. [16] to develop our maturity model because it is based on DSR and therefore provides a methodological foundation very suitable for application in our research approach. Furthermore, Becker et al. provide a stringent and consistent development process according to the DSR guidelines of Hevner et al. [17].

Becker et al. [16] argue that maturity models are artifacts that serve to solve the problem of appreciating capacity and obtain improvement measures. According to [19] design science allows you to create artifacts such as constructs, models, methods, and

instantiations that help improve problem-solving capabilities. Thus, the authors state that design science research is appropriate for the development of maturity models.

In the same study [16], the author proposes a procedure for the development of maturity models composed of eight steps. All steps should be documented. As depicted in the procedure model in Fig. 2 the first steps focus on the problem identification (step 1). In this step the research problem is identified and detailed, the practical relevance of the problem is specified and the value of the artifact is justified. This step is followed by the comparison with existing maturity models (step 2). This second step is based on the problem identification of the first step and analysis of existing maturity model in the Information Security Management domain, which leads to the identification of weaknesses in these models. We conducted a literature analysis, which was based on an extensive online search to find existing maturity models focused on the Information Security Management domain. Thus, the analysis of the maturity models was performed according to their functionality, as well as, their capability to address the ISO/IEC 27001 requirements.




Fig. 2. Procedure model of the research approach (adopted from Becker et al. [16])

The next step deals with the determination of the research strategy (step 3) outlined in this section of the paper. This is followed by the iterative maturity model development (step 4). In this step, we used model adoption techniques, such as, configuration, instantiation, aggregation, specialization and analogy [18] to incorporate the ISO/IEC 27001 in the maturity model. This allowed us to create a rigorous maturity model regarding both the structure and content. In the last step, evaluation (step 5), we combined the steps of Becker et al. [16], conception of transfer and evaluation, implementation of transfer media, and evaluation. All steps will be conducted, but to match the structure of this paper we made this change.

4 Problem Analysis

In order to provide a consistent and precise problem definition, we gathered the ISMS process requirements from ISO/IEC 27001. According to the ISO/IEC 27001, the activities for ISMS Processes can be summarized as follows:

- A1: Establish the ISMS – “Establish ISMS policy, objectives, processes and procedures relevant to managing risk and improving information security to deliver results in accordance with an organization’s overall policies and objectives.” [2];
- A2: Implement and operate the ISMS – “Implement and operate the ISMS policy, controls, processes and procedures.” [2];
- A3: Monitor and Review the ISMS – “Assess and, where applicable, measure process performance against ISMS policy, objectives and practical experience and report the results to management for review.” [2];
- A4: Maintain and Improve the ISMS – “Take corrective and preventive actions, based on the results of the internal ISMS audit and management review or other relevant information, to achieve continual improvement of the ISMS.” [2].

Table 1. ISO/IEC 27001 activities reference matrix fit assessment

Maturity model	A1	A2	A3	A4	Σ
O-ISM3	2	3	4	4	13
SSE-CMM	2	4	4	2	12
ISF MM	2	2	3	3	10
COBIT 5	4	2	4	2	12
ONG C2M2	3	2	2	3	10
BSIMM	3	4	4	4	15
Average	2,6	2,8	3,5	3	12

These are the activities that the ISMS process must perform in order to be in line with the recommendations of the ISO/IEC 27001. The activities are used as a reference baseline to assess the appropriateness of several existing Information Security Management Maturity Models. Based on the results of the literature review we conducted within the Information Security Management domain, we identified several papers dealing with maturity models. We selected maturity models that used different methodological approaches. Then, each maturity model was analyzed according to the degree to which they cover and fit to the previously defined reference baseline. Each maturity model was ranked for every requirement according to the degree of matching, using a Likert-scale, from 1 (very low) to 5 (very high). After this analysis, we concluded that only six maturity models scored an aggregate of at least 10 points according to the defined ISO/IEC 27001 activities baseline: (1) Open Information Security Management Maturity Model (O-ISM3) [6]; (2) Systems Security Engineering – Capability Maturity Model (SSE-CMM) [7]; (3) ISF Maturity Model Accelerator (ISF MM) [8]; (4) Control Objectives for Information and Related Technologies - Version 5 (COBIT 5) [9]; (5) Cyber Security Capability Maturity Model (C2M2) [10], and (6) Building Security in Maturity Model

(BSIMM) [11]. Table 1 presents the assessment results of the above as the most significant identified maturity model in detail. Based on this set an average total score of 12 was achieved (maximum score 20).

5 Maturity Model Design

In accordance to the maturity model development approach of Becker et al. [16] a new maturity model has to be developed, if no existing or the advancement of an existing one is capable of addressing the identified problem. So, based on the findings of our analysis there is no maturity model which satisfactorily fulfill the entire ISO/IEC 27001 activities baseline. Therefore, we will develop a new maturity model. The newly developed maturity model presented in Table 2 adopts established structural elements, domains and functions of the best practice in ISO/IEC 27001. As detailed within the research methodology, we applied an iterative process for the maturity model development. In total we needed two iterations which can be detailed as follows:

First iteration: As a first step, we defined the characteristics and structure of the maturity model. We started by proposing five maturity levels, Initial, Managed, Defined, Quantitatively Managed, and Optimizing. These maturity levels can be found in various established maturity models, such as, CMMI [4]. In this initial iteration, we focused in just a part of the ISO/IEC 27001 ISMS process namely the Plan step. For each criterion of the maturity model we modeled what was the manifestation of that criterion at the different maturity levels.

Second Iteration: In the second iteration we completely overhauled the definition of the maturity levels by proposing five new maturity levels, Initial, Planning, Implementation, Monitoring, and Improvement. These maturity levels are based on the PDCA cycle used within the ISO/IEC 27001 as depicted in Fig. 1. Table 3 details the activities on which our maturity model is based, along with a mapping to the ISO/IEC 27001 ones they were derived from. This made it easier for a user accustomed with the ISO/IEC 27001 to understand the maturity model and make a connection between what was being asked in each assessment criterion and the requirements specified in the ISO/IEC 27001, which resulted in an easily understandable maturity model that is presented in Table 2. Finally, this leads to the following maturity levels: (Level 1) Initial Stage; (Level 2) Planning Stage; (Level 3) Implementation Stage; (Level 4) Monitoring Stage; (Level 5) Improvement Stage.

To improve from level X to level $X + 1$, the organization must comply with all the criteria from level X , which makes this maturity model follow a “stages” approach. What an organization can expect from progressing through the maturity levels is that their ISMS process will become increasingly managed, defined and optimized.

Table 2. ISMS maturity model

Maternity level	Assessment criterion
Level 1: initial	<i>No criteria</i>
Level 2: planning	2.1 - Define scope and boundaries of the ISMS
	2.2 - Define ISMS policy
	2.3 - Define risk assessment approach
	2.4 - Perform risk identification
	2.5 - Perform risk analysis and evaluation
	2.6 - Define risk treatment options
	2.7 - Define risk treatment control objectives and controls
	2.8 - Obtain management approval for residual risks
	2.9 - Obtain management authorization to implement and operate the ISMS
	2.10 - Prepare a statement of applicability
Level 3: implementation	3.1 - Formulate risk treatment plan
	3.2 - Implement risk treatment plan
	3.3 - Implement controls to meet control objectives
	3.4 - Define effectiveness measurement procedure of the selected controls
	3.5 - Implement training and awareness programmes
	3.6 - Manage operation of the ISMS
	3.7 - Manage resources for the ISMS
	3.8 - Implement procedures and controls for detection and response to security events
Level 4: monitoring	4.1 - Execute monitoring and reviewing procedures and other controls
	4.2 - Undertake regular reviews of the effectiveness of the ISMS
	4.3 - Measure the effectiveness of controls
	4.4 - Review risk assessments
	4.5 - Review residual risks
	4.6 - Review identified acceptable levels of risks
	4.7 - Conduct internal ISMS audits
	4.8 - Undertake management review of the ISMS
	4.9 - Update security plans
	4.10 - Record actions and events
Level 5: improvement	5.1 - Implement the identified improvements in the ISMS
	5.2 - Take appropriate corrective and preventive actions
	5.3 - Communicate actions and improvements to all interested parties
	5.4 - Ensure that improvements achieve their objectives

6 Maturity Model Evaluation

The evaluation step is a main element of DSR. It is necessary to show the “utility, quality, and efficacy of a design artifact” [19]. To be compliant with these requirements we evaluated the ISMS Maturity Model by using a multi-perspective approach which consists of three stages: (1) Evaluation of the mapping using the Wand and Weber

Table 3. Mapping of the ISMS maturity model and ISO/IEC 27001 requirements, and the resulting evaluation using the Wand and Weber (W&W) ontological deficiencies.

PCDA cycle	ISMS maturity model activities	ISO/IEC 27001 requirements	W&W ontological deficiencies
Plan	<i>Maturity level: planning</i>		
	Define scope and boundaries of the ISMS	4.2.1 – (a)	Complete
	Define ISMS Policy	4.2.1 – (b)	Complete
	Define risk assessment approach	4.2.1 – (c)	Complete
	Risk identification	4.2.1 – (d)	Complete
	Risk analysis and evaluation	4.2.1 – (e)	Complete
	Risk treatment options	4.2.1 – (f)	Complete
	Risk treatment control objectives and controls	4.2.1 – (g)	Complete
	Obtain management approval for residual risks	4.2.1 – (h)	Complete
	Obtain management authorization to implement and operate the ISMS	4.2.1 – (i)	Complete
Prepare a statement of applicability	4.2.1 – (j)	Complete	
Do	<i>Maturity level: implementation</i>		
	Formulate risk treatment plan	4.2.2 - (a)	Complete
	Implement risk treatment plan	4.2.2 - (b)	Complete
	Implement controls to meet control objectives	4.2.2 - (c)	Complete
	Define effectiveness measurement procedure of the selected controls	4.2.2 - (d)	Complete
	Implement training and awareness programmes	4.2.2 - (e)	Complete
	Manage operation of the ISMS	4.2.2 - (f)	Complete
	Manage resources for the ISMS	4.2.2 - (g)	Complete
Implement procedures and controls for detection and response to security events.	4.2.2 - (h)	Complete	
Check	<i>Maturity level: monitoring</i>		
	Execute monitoring and reviewing procedures and other controls	4.2.3 – (a)	Complete
	Undertake regular reviews of the effectiveness of the ISMS	4.2.3 – (b)	Complete
	Measure the effectiveness of controls	4.2.3 – (c)	Complete
	Review risk assessments	4.2.3 – (d)	Overload
	Review residual risks	4.2.3 – (d)	Overload
	Review identified acceptable levels of risks	4.2.3 – (d)	Overload
	Conduct internal ISMS audits	4.2.3 – (e)	Complete
	Undertake management review of the ISMS	4.2.3 – (f)	Complete
	Update Security Plans	4.2.3 – (g)	Complete
Record Actions and Events	4.2.3 – (h)	Complete	
Act	<i>Maturity level: improvement</i>		
	Implement the identified improvements in the ISMS	4.2.4 – (a)	Complete
	Take appropriate corrective and preventive actions	4.2.4 – (b)	Complete
	Communication	4.2.4 – (c)	Complete
	Ensure that improvements achieve their objectives	4.2.4 – (d)	Complete

Ontological Deficiencies; (2) an assessment of the fit of the ISMS Maturity Model against the ISO/IEC 27001 requirements used to compare existing ISMS maturity models in Sect. 4; and (3) assess five organizations using the ISMS Maturity Model.

To evaluate the mapping between our maturity model and ISO/IEC 27001, regarding completeness and clarity, we performed an analysis according to the Wand and Weber method [20]. Wand and Weber define an ontological evaluation of the grammars method, where two sets of concepts are compared in order to identify four ontological deficiencies, as depicted in Fig. 3:

- **Incompleteness** - Can every element in the first set be mapped to an element in the second set? If there is not a total mapping, it is considered incomplete;
- **Redundancy** - Are there elements in the first set mapped to more than one element in the second set? If so, the mapping is considered redundant;
- **Excess** - Is every element from the second set mapped to an element in the second set? The mapping is considered excessive if there are elements from the second set without a relationship;
- **Overload** - Is every element of the second set mapped to only one element in the first set? The mapping is considered overloaded if any element in the second set has more than one mapping to the first set.




Fig. 3. Wand and weber ontological deficiencies [20]

The ontological evaluation of the mapping between the ISMS Maturity Model and ISO/IEC 27001 chapters (see Fig. 1) and requirements is detailed in Table 3. A first observation is that the mapping is complete, since every proposed activity can be mapped to an ISO/IEC 27001 requirement. As for the other attributes, there is no redundancy and excess. However, regarding overload, the ISO/IEC 27001 “4.2.3 - (d)” requirement was overloaded as in our understanding it defines a requirement for three different activities. As a result, we created three different assessment criteria for this requirement. Finally, the ISMS Maturity Model covers all the requirements detailed in Sect. 4, which means that the total score using the same scale is 20.

Following the first two evaluation steps, we assessed five real organizations by following an assessment method, anonymized due to consent issues. Organization Alpha is the public institute responsible for promoting and developing administrative modernization in its country. Its operation is in three axes: customer service, digital transformation and simplification. Organization Beta is part of the business sector in its country government that produces and supplies goods and services that require high security standards, namely: coins, banknotes, and documents, such as, citizen’s card and

passports. Organization Gamma is a public higher education institution that has approximately 11.500 students being the largest school of engineering, science and technology in its country. Organization Delta is a public institution for scientific and technological research and development whose purpose is to contribute to the creation, development and diffusion of research in fields related to civil engineering. Organization Omega is a private organization which focus on software development and maintenance providing services all over the globe with various offices in Europe.

For each of these five organizations we took the role of assessors, assessed the organization collecting objective evidence for the assessment criteria defined in the maturity model. Then, the results were analyzed which resulted in the assessment results depicted in Table 4. In this table, “Y” stands for criterion satisfied, an empty cell stands for criterion not satisfied, and the last columns shows the final maturity level for each of the assessed organizations.

Table 4. ISMS maturity assessment results

	Criterion	2.1	2.2	2.3	2.4	2.5	2.6	2.7	2.8	2.9	2.10	3.1	3.2	3.3	3.4	3.5	3.6	3.7	3.8
Organization	ALPHA	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
	BETA	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
	GAMMA	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y					Y	Y	Y	
	DELTA	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
	OMEGA	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

	Organization	4.1	4.2	4.3	4.4	4.5	4.6	4.7	4.8	4.9	4.10	5.1	5.2	5.3	5.4	Maturity Level
Organization	ALPHA	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	5
	BETA						Y	Y	Y							3
	GAMMA			Y					Y							2
	DELTA	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y					4
	OMEGA	Y	Y	Y		Y						Y	Y			3

In order to achieve a certain maturity level, the organization must comply with all the criteria for that specific level and the levels below, which means that an organization at maturity level 3 complies with all the criteria for maturity levels 2 and 3.

As can be perceived from Table 4, we were able to assess each of the assessment criteria, which in turn allowed us to determine the ISMS maturity level for each of the five organizations. From our analysis, the assessment results shown that the maturity model correctly determined the maturity levels and these in fact correspond to our perception of the maturity of the ISMS implemented in the organization. These results were then used by the organizations to create improvement plans specially tailored to their organizational context.

7 Conclusions

The aim of this paper is to detail the development of a maturity model for the ISMS process based on the ISO/IEC 27001 standard. The latter can serve as a governance instrument that could be used by the Information Security Management function to analyze and evaluate the current strengths and weaknesses of the ISMS process.

However, the model is not restricted to analytical purposes only. It can also be used to derive a roadmap towards an evolutionary improvement of the Information Security Management function regarding its capabilities and its effectiveness and efficiency.

The first part of the paper elaborates the ISMS activities requirements which were used as a reference baseline to investigate whether existing maturity models are capable of holistically assessing an ISMS process (RQ1). The findings revealed that existing maturity models cover the entire reference baseline insufficiently, since they only selectively address the activities. Hence, no existing maturity model is able to solve the identified problem. Finally, we decided to design a new maturity model in consistency to the defined research strategy.

In the second part of the paper, we described the development of a maturity model for ISMS, including the model itself as well as its evaluation to address the second research question (RQ2). The developed model is based on existing maturity model structures and inherits concepts and methodologies of the ISO/IEC 27001. The researchers took care during the development to provide a consumable research result. Moreover, the ISMS maturity model benefited from the multi-perspective evaluation approach by further advancements.

Naturally, the applied research approach comes along with certain limitations. This paper presents the assessment results for five organizations using the ISMS Maturity Model. However, in order to extend usefulness of the maturity model, as well as, provide additional validation scenarios and further improve the research aspect, we suggest evaluating (and refining) the ISMS maturity model within different industry sectors, this would lead to a more generic ISMS maturity model and would enable cross-industry benchmarking.

Acknowledgements. This work was supported by national funds through Fundação para a Ciência e a Tecnologia (FCT) with reference UID/CEC/50021/2013.

References

1. Dubois, E., Heymans, P., Mayer, N., Matulevicius, R.: A systematic approach to define the domain of information system security risk management. In: Nurcan, S., Salinesi, C., Souveyet, C., Ralyté, J. (eds.) *Intentional Perspectives on Information Systems Engineering*, pp. 289–306. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-12544-7_16
2. ISO/IEC 27001:2013, Information technology - Security techniques - Information security management systems – Requirements (2013)
3. Miller, H., Murphy, R.: *Secure cyberspace: answering the call for intelligent action*. IT Professional (2009)
4. CMMI Product Team: *CMMI for Development, Version 1.3*, Carnegie Mellon Univ., no. November, p. 482 (2010)
5. Mettler, T.: *A design science research perspective on maturity models in information systems*. Institute of Information Management, University of St. Gallen, St. Gallen (2009)
6. The Open Group: *Open Information Security Management Maturity Model (O-ISM3)* (2011)
7. Carnegie-Mellon-University: *Systems Security Engineering Capability Maturity Model (SSE-CMM) - Model Description Document. Version 3.0* (2003)

8. ISF: Time to grow using maturity models to create and protect value, in Information Security Forum ISF (2014)
9. IT Governance Institute: COBIT 5 – A business Framework for the Governance and Management of Enterprise IT (2012)
10. Department of Energy, U.S. Department of Homeland Security, Cybersecurity Capability Maturity Model (C2M2 v1.1) (2014)
11. McGraw, G., Miguez, S., West, J.: Building Security in Maturity Model (BSIMM) Version 8 (2015)
12. Mettler, T., Rohner, P., Winter, R.: Towards a classification of maturity models in information systems. In: D’Atri, A., De Marco, M., Braccini, A., Cabiddu, F. (eds.) Management of the Interconnected World. Physica-Verlag, Heidelberg (2010). https://doi.org/10.1007/978-3-7908-2404-9_39
13. De Bruin, T., Freeze, R., Kaulkarni, U., Rosemann, M.: Understanding the main phases of developing a maturity assessment model. In: Proceedings of the Australasian Conference on Information Systems (ACIS) (2005)
14. Poeppelbuss, J., Niehaves, B., Simons, A., Becker, J.: Maturity models in information systems research: literature search and analysis. In: Communications of the Association for Information Systems, vol. 29 (2011)
15. van Loon, H.: Process Assessment and Improvement: A Practical Guide. Springer, New York (2015)
16. Becker, J., Knackstedt, R., Pöppelbuß, J.: Developing maturity models for IT management: a procedure model and its application. *Bus. Inf. Syst. Eng.* **3**, 213–222 (2009)
17. Hevner, A., Ram, S., March, S., Park, J.: Design science in information systems research. *MISQ* **28**, 75–105 (2004)
18. Vom Brocke, J.: Design principles for reference modeling-reusing information models by means of aggregation, specialization, instantiation, and analogy. In: Fettke, P., Loos, P. (eds.) Reference Modeling for Business Systems Analysis. Idea Group Inc., Hershey (2007)
19. Hevner, A., Chatterjee, S.: Design Research in Information Systems: Theory and Practice. Springer, Heidelberg (2010). <https://doi.org/10.1007/978-1-4419-5653-8>
20. Wand, Y., Weber, R.: On the ontological expressiveness of information systems analysis and design grammars. *Inf. Syst. J.* **3**(4), 217–237 (1993)



Repairing Outlier Behaviour in Event Logs

Mohammadreza Fani Sani¹(✉), Sebastiaan J. van Zelst²,
and Wil M. P. van der Aalst¹

¹ Process and Data Science Chair, RWTH Aachen University,
Ahornstraße 55, Erweiterungsbau E2, 52056 Aachen, Germany
{fanisani,wvdaalst}@pads.rwth-aachen.de

² Department of Mathematics and Computer Science,
Eindhoven University of Technology,
P.O. Box 513, 5600 MB Eindhoven, The Netherlands
s.j.v.zelst@tue.nl

Abstract. One of the main challenges in applying process mining on real event data, is the presence of noise and rare behaviour. Applying process mining algorithms directly on raw event data typically results in complex, incomprehensible, and, in some cases, even inaccurate analyses. As a result, correct and/or important behaviour may be concealed. In this paper, we propose an event data repair method, that tries to detect and repair outlier behaviour within the given event data. We propose a probabilistic method that is based on the occurrence frequency of activities in specific contexts. Our approach allows for removal of infrequent behaviour, which enables us to obtain a more global view of the process. The proposed method has been implemented in both the ProM- and the RapidProM framework. Using these implementations, we conduct a collection of experiments that show that we are able to detect and modify most types of outlier behaviour in the event data. Our evaluation clearly demonstrates that we are able to help to improve process mining discovery results by repairing event logs upfront.

Keywords: Process mining · Data cleansing · Log repair
Event log preprocessing · Conditional probability · Outlier detection

1 Introduction

Process Mining bridges the gap between traditional data analysis techniques like data mining and business process management analysis [1]. It aims to discover, monitor, and enhance processes by extracting knowledge from event data, also referred to as *event logs*, readily available in most modern information systems [2]. In general, we identify three main branches of process mining being *process discovery*, *conformance checking*, and *process enhancement*. In process discovery, we try to discover a process model that accurately describes the underlying

process captured within the event data. In conformance checking we try to assess to what degree a given process model (possibly the result of a process discovery algorithm) and event data conform to one another. Finally, process enhancement aims at improving the view on a process by improving its corresponding model based on the related event data.

Most process mining algorithms, in any of these branches, work under the assumption that behaviour related to the execution of the underlying process is stored correctly within the event log. Moreover, completeness of the behaviour, i.e., each instance of the process as stored in the event log is already finished, is assumed as well. However, real event data often contains noise, i.e., behaviour that is/should not be part of the process. Furthermore, it often contains data related to infrequent behaviour, i.e., behaviour that is rather rare due to the handling of exceptional cases. The presence of such behaviour, which we subsequently refer to as *outlier*, makes most process mining algorithms return complex, incomprehensible or even inaccurate results. To reduce these negative effects, process mining projects often comprise of a pre-processing phase in which one tries to detect and remove traces that contain such undesired behaviour. This cleaning phase is usually performed manually and is therefore rather costly and time consuming.

Despite the negative impacts of the presence of noisy, incomplete and infrequent behaviour, little research has been done towards automated data cleansing techniques that improve process mining results. Recently, research has been performed aiming at *filtering* out traces that contain outlier behaviour from an event log [3,4]. Even though both techniques show improvements in process mining results, in particular w.r.t. process discovery, only a little fragment of outlier behaviour within a trace of event data leads to ignoring the trace as a whole. This potentially leads to a distortion of the general distribution of common behaviour of the process, yielding potentially inaccurate or even wrong process mining results.

Therefore, we propose a general purpose *event log repairing technique* that, given event potentially containing outlier behaviour, identifies and modifies such behaviour in order to obtain a more reliable input for all possible process mining algorithms. In particular, we use a probabilistic method to detect outlier behaviour according to the *context* of a process i.e., fragments of activity sequences that occur before and after the potential outlier behaviour. After outlier identification, the corresponding behaviour is replaced with another fragment of behaviour that is more probable to occur according to the context in which the outlier behaviour occurs.

Using the ProM [5] based extension of RapidMiner, i.e., RapidProM [6], we study the effectiveness of our approach, using both synthetic and real event data. The obtained results show that our approach adequately identifies and repairs outlier behaviour and as a consequence increases the overall quality of process model discovery results. Additionally, we show that our method achieves better results compared to one of the best performing existing filtering techniques in the process mining domain.

The remainder of this paper is structured as follows. Section 2 motivates the need for data cleansing and repair methods in context of process mining. In Sect. 3, we discuss related work. We present our proposed outlier repair method in Sect. 4. Evaluation details and corresponding results are given in Sect. 5. Finally, Sect. 6 concludes the paper and presents directions for future work.

2 Motivation

Real event logs often contain noise/anomalous behaviour. The presence of such behaviour causes many problems for process mining algorithms. In particular, most process discovery algorithms incorporate all behaviour in event logs as much as possible. As a result, most outlier behaviour is incorporated as well which decreases the overall accuracy of the discovered model. Therefore, it is essential to accurately pre-process event data. In fact, according to the process mining manifesto [7], cleaning event logs is one of the main challenges in the *process mining* field.

It is possible to define *outlier behaviour* in a variety of ways. For example, in [2], *infrequent behaviour* is considered as outlier behaviour. In turn, infrequent behaviour relates to rare behaviour that is supposed to be part of the process, yet severely hampers the feasibility of process mining results. In practice, behaviour that is not part of the process, i.e. caused by logging exceptions, faulty execution of the process, is infrequent by its sheer nature. So, in this paper we define both infrequent behaviour that is part of the process and (infrequent) faulty execution behaviour, i.e. noise, as outlier behaviour.

An example event log with some outlier behaviour is shown in Table 1. In this event log there are 74 events belong to 11 traces. Except for the first variant (unique behaviour of a process instance) each variant has just one corresponding trace, note that this is customary in some application domains, e.g. medical treatment process [8]. The first three traces contain no outlier behaviour. However, the next seven variants have different types of outlier behaviour. For example, in the fourth and fifth rows the activities “*h*” and “*a*” are missing. Some process discovery algorithms like the Alpha miner [9] are sensitive to such outlier behaviour and yield inferior process discovery results when applied directly to such event logs. Other process discovery algorithms like the Inductive Miner [10], have embedded filtering mechanisms to deal with some types of outliers. The results of applying various process discovery algorithms on this event log are shown in Fig. 1. If we apply filtering method of [4] on this event log, variants 1 – 5 are retained. A resulting process model using these traces combined with the Inductive Miner is shown in Fig. 1e. As it is shown in this figure, all mentioned process discovery algorithms have problem to discover an accurate model from the given event log. However, if we first repair the event log and then apply the Alpha miner (or any other mentioned process discovery methods), we obtain an accurate process model, i.e. as presented in Fig. 1f. It is because there is no outlier behaviour in the repaired event log and the resulted process model is more accurate and understandable.

Table 1. Event log with 11 traces and 10 different trace-variants.

Variant	Frequency
$\langle a, b, c, d, e, f, h \rangle$	2
$\langle a, b, d, c, e, f, h \rangle$	1
$\langle a, b, c, d, e, g, h \rangle$	1
$\langle a, b, d, c, e, g \rangle$	1
$\langle b, d, c, e, g, h \rangle$	1
$\langle a, b, c, d, e \rangle$	1
$\langle a, b, c, d, g, h \rangle$	1
$\langle a, b, b, c, d, e, f, h \rangle$	1
$\langle a, b, c, d, g, h \rangle$	1
$\langle a, b, d, c, a, e, g, f, h \rangle$	1

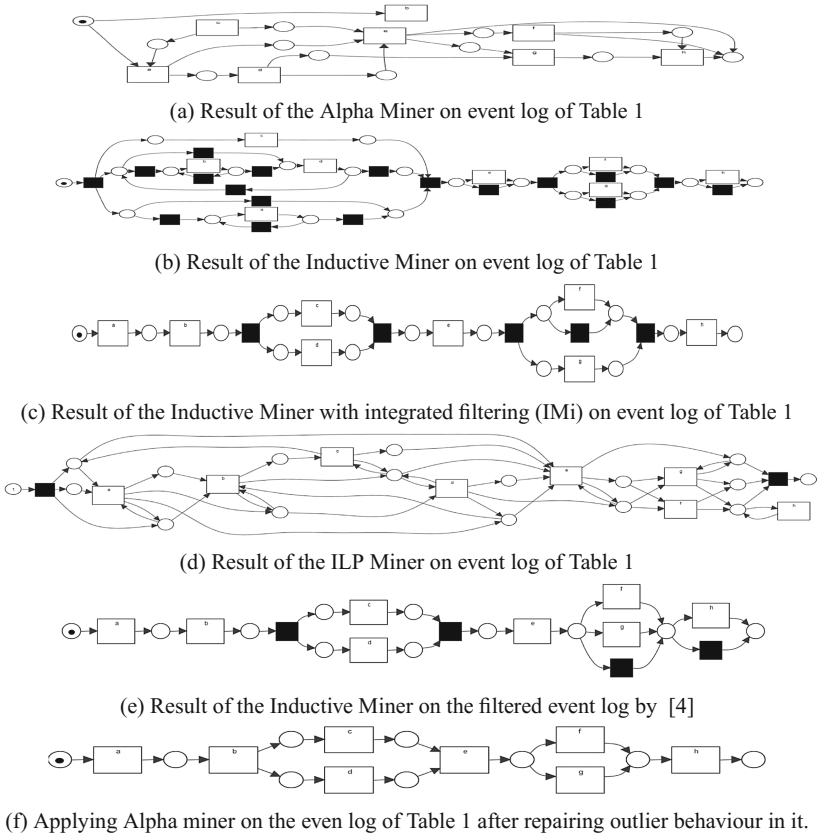


Fig. 1. Resulting process models of applying different process discovery techniques (with and without filtering/repair) on the event log that is presented in Table 1

So, it seems that we need an approach to overcome outlier behaviour in event data. A naive approach to solve data quality related issues is to remove traces that seem to contain outlier behaviour [3,4]. However, for many businesses, all process instances in an event log are valuable and ignoring them potentially jeopardizes the trustworthiness of the analysis performed. For example, in a patient treatment in a hospital, recorded over several years, it is undesirable to remove all process related records of a patient just because there exists a small portion of wrongly logged behaviour. In such scenarios, after detecting outlier, it is more desirable to repair that behaviour. Note that, if there is no noise in an event log, by repairing infrequent behaviour in it and removing too detailed patterns, we are able to alter infrequent behaviour into more frequent behaviour which allows us to discover more general views on the process. Therefore, by repairing, rather than removing outlier behaviour we believe that quality of discovered process models improves.

3 Related Work

Some process mining algorithms are designed to/have been extended to be able to handle outliers as well [11–14]. However, these filtering techniques are tailored towards the internal working of the corresponding algorithms and cannot be used for general purpose event log cleansing. Additionally, they typically focus on a specific type of noise, e.g. incompleteness.

Most of filtering techniques present in commercial tools are based on just the frequency of activities and variants. However, the presence of parallelism and loops often hampers the applicability of such filtering techniques. There are also some basic filtering plug-ins developed in ProM [5] based on activity frequencies and users inputs.

Outlier detection for general temporal data is addressed in some research, e.g. in [15] a survey on different methods of detecting outliers in sequential data is presented. Also, there are some related techniques that specifically proposed for process mining domain. In [16,17] the authors propose filtering techniques that use additional information such as training event data or a reference process model. In reality, providing a sufficiently complete set of training traces or having a reference process model is impractical. Recently, some general purpose filtering techniques are proposed in the process mining domain. In [3] by constructing an Anomaly Free Automaton (AFA) based on the whole event log and a given threshold, all non-fitting behaviour, w.r.t. the AFA, is removed from the event log. In [4], we propose a filtering method that detects outliers based on conditional probabilities of subsequences and their possible following activities. In [18], an adjustable on-line filtering method is proposed that detect outlier behaviour for streaming event logs that also works based on conditional probabilities.

All aforementioned methods, after detecting outlier behaviour, try to remove such behaviour. As motivated in Sect. 2, modifying outlier behaviour or repairing it is more valuable than just removing it. There is some research [19,20] that

tries to repair process models based on event logs. Moreover, [21] uses process model to repair event logs. However, as we aim to design a general purpose repairing method, we assume there does not exist a process model as an input of our proposed algorithm.

4 Repairing Outliers in Event Data

In this section we present our outlier repair method. Prior to this, we briefly introduce basic process mining terminology and notations that ease readability of the paper.

4.1 Terminology and Notation

Given a set X , a multiset M over X is a function $M: X \rightarrow \mathbb{N}_{\geq 0}$, i.e. it allows certain elements of X to appear multiple times. We write a multiset as $M = [e_1^{k_1}, e_2^{k_2}, \dots, e_n^{k_n}]$, where for $1 \leq i \leq n$ we have $M(e_i) = k_i$ with $k_i \in \mathbb{N}_{> 0}$. If $k_i = 1$, we omit its superscript, and if for some $e \in X$ we have $M(e) = 0$, we omit it from the multiset notation. Also, $M = []$ denotes an empty multiset, i.e. $\forall e \in X, M(e) = 0$. We let $\overline{M} = \{e \in X \mid M(e) > 0\}$, i.e. $\overline{M} \subseteq X$. The set of all possible multisets over a set X is written as $\mathcal{M}(X)$.

Let X^* denote the set of all possible sequences over a set X . A finite sequence σ of length n over X is a function $\sigma: \{1, 2, \dots, n\} \rightarrow X$, alternatively written as $\sigma = \langle x_1, x_2, \dots, x_n \rangle$ where $x_i = \sigma(i)$ for $1 \leq i \leq n$. The empty sequence is written as ϵ . Concatenation of sequences σ and σ' is written as $\sigma \cdot \sigma'$. We let function $hd: X^* \times \mathbb{N}_{\geq 0} \rightarrow X^*$, represents the “head” of a sequence, i.e., given a sequence $\sigma \in X^*$ and $k \leq |\sigma|$, $hd(\sigma, k) = \langle x_1, x_2, \dots, x_k \rangle$, i.e., the sequence of the first k elements of σ . In case $k = 0$ we have $hd(\sigma, 0) = \epsilon$. Symmetrically, $tl: X^* \times \mathbb{N}_{\geq 0} \rightarrow X^*$ represents the “tail” of a sequence and is defined as $tl(\sigma, k) = \langle x_{n-k+1}, x_{n-k+2}, \dots, x_n \rangle$, i.e., the sequence of the last k elements of σ , with, again, $tl(\sigma, 0) = \epsilon$. Sequence σ' is a subsequence of sequence σ , which we denote as $\sigma' \in \sigma$, if and only if $\exists \sigma_1, \sigma_2 \in X^* (\sigma = \sigma_1 \cdot \sigma' \cdot \sigma_2)$. Let $\sigma, \sigma' \in X^*$. We define the frequency of occurrence of σ' in σ by $freq: X^* \times X^* \rightarrow \mathbb{N}_{\geq 0}$ where $freq(\sigma', \sigma) = |\{1 \leq i \leq |\sigma| \mid \sigma'_1 = \sigma_i, \dots, \sigma'_{|\sigma'|} = \sigma_{i+|\sigma'|}\}|$. Given an example event log that presented in Table 1, $freq(\langle b \rangle, \langle a, b, b, c, d, e, f, h \rangle) = 2$ and $freq(\langle b, d \rangle, \langle a, b, d, c, e, g \rangle) = 1$, etc.

Event logs describe sequences of executed business process activities, typically in context of some cases (or process instances), e.g., a customer or an order-id. The execution of an activity in context of a case is referred to an *event*. A sequence of events for a specific case is also referred to a *trace*. Thus, it is possible that multiple traces describe the same sequence of activities, yet, since events are unique, each trace itself contains different events. An example event log, adopted from [2], is presented in Table 2.

Consider the events related to *Case-id* value 1. Sara registers a request, after which Ali examines it thoroughly. William checks the ticket and checks resources. Ava sends the request to manager and Fatima accept the request. Finally, Anna

emails the decision to the client. The example trace is written as $\langle a, b, c, d, e, f, h \rangle$ (using short-hand activity names). In the context of this paper, we formally define event logs as a multiset of sequences of activities (e.g., Table 1).

Table 2. Fragment of a fictional event log (each line corresponds to an event).

Case-id	Activity	Resource	Time-stamp
...
1	<i>register request (a)</i>	<i>Sara</i>	<i>2017-04-08:08.10</i>
1	<i>examine thoroughly (b)</i>	<i>Ali</i>	<i>2017-04-08:09.17</i>
2	<i>register request (a)</i>	<i>Sara</i>	<i>2017-04-08:10.14</i>
1	<i>check resources (c)</i>	<i>William</i>	<i>2017-04-08:10.23</i>
1	<i>check ticket (d)</i>	<i>William</i>	<i>2017-04-08:10.53</i>
2	<i>check resources (b)</i>	<i>Ali</i>	<i>2017-04-08:11.13</i>
1	<i>Send to manager (e)</i>	<i>Ava</i>	<i>2017-04-08:13.09</i>
1	<i>accept request (f)</i>	<i>Fatima</i>	<i>2017-04-08:16.05</i>
1	<i>mail decision (h)</i>	<i>Anna</i>	<i>2017-04-08:16.18</i>
...

Definition 1 (Event Log). Let \mathcal{A} be a set of activities. An event log is a multiset of sequences over \mathcal{A} , i.e. $L \in \mathcal{M}(\mathcal{A}^*)$.

Observe that each $\sigma \in \bar{L}$ describes a *trace-variant* whereas $L(\sigma)$ describes how many traces of the form σ are present within the event log.

4.2 Repairing Event Logs Using Control-Flow Oriented Contexts

Here, we present our methodology to identify and repair outlier behaviour in an event log. We first present our repair method by means of an example after which we formalize and discuss it in more detail. We first present two central control-flow oriented concepts which form the basis of the repair method. Given a trace, a context is defined as the surrounding behaviour of a certain sequence of activities. For example, in trace $\langle a, b, c, d, e, f, h \rangle$, $\langle a, b \rangle$ and $\langle e \rangle$ are surrounding $\langle c, d \rangle$, hence, the pair $(\langle a, b \rangle, \langle e \rangle)$ is a context of $\langle c, d \rangle$.

Definition 2 (Context, Covering). Let $\sigma, \sigma' \in X^*$. We define the context of σ' w.r.t. σ as a function $con : X^* \times X^* \times \mathbb{N}_{\geq 0} \times \mathbb{N}_{\geq 0} \rightarrow \mathcal{P}(X^* \times X^*)$ where $con(\sigma', \sigma, l, r) = \{(\sigma_1, \sigma_2) \in X^* \times X^* \mid \sigma_1 \cdot \sigma' \cdot \sigma_2 \in \sigma \wedge |\sigma_1| = l \wedge |\sigma_2| = r\}$.

Furthermore, let $\sigma'_1, \sigma'_2 \in X^*$, $cov : X^* \times X^* \times X^* \rightarrow \mathcal{P}(X^*)$ is a function that returns all subsequences in a trace that occur within a given context, i.e. $cov(\sigma, \sigma'_1, \sigma'_2) = \{\sigma' \in X^* \mid (\sigma'_1, \sigma'_2) \in con(\sigma', \sigma, |\sigma'_1|, |\sigma'_2|)\}$.

A context describes, given a subsequence, its two neighbouring subsequences of length l and r respectively, i.e. σ'_1 is the left neighbour with length l and σ'_2 is the right neighbor with length r . For example, if $\sigma = \langle a, b, c, d, e, f, h \rangle$, we have $\text{con}(\langle c \rangle, \sigma, 1, 2) = (\langle b \rangle, \langle d, e \rangle)$. Note that l and r may differ, and, may even have value equal to 0 which implies a neighboring subsequence ϵ . Furthermore, $\text{cov}(\sigma, \langle a, b \rangle, \langle e, f \rangle) = \{\langle c, d \rangle\}$ and $\text{cov}(\sigma, \langle b \rangle, \langle d \rangle) = \{\langle c \rangle\}$.

We aim to detect and replace outlier behaviour based on the occurrence probability of a subsequence among a specific context. If this probability is lower than a given threshold, we consider that behaviour as outlier. To obtain these probabilities we define *covering probability* as follows.

Definition 3 (Covering Probability). *Given $\sigma', \sigma'_1, \sigma'_2 \in X^*$, maximum subsequences' length K and a multiset of sequences $L \in \mathcal{M}(X^*)$. We define $CP: X^* \times X^* \times X^* \times \mathbb{N}_{>0} \times \mathcal{M}(X^*) \rightarrow [0, 1]$ as the empirical conditional probability of σ' being covered by σ_1 and σ_2 in event log L , i.e.*

$$CP(\sigma', \sigma'_1, \sigma'_2, K, L) = \begin{cases} \frac{\sum_{\sigma \in L} (L(\sigma) \text{freq}(\sigma'_1 \cdot \sigma' \cdot \sigma'_2, \sigma))}{\sum_{\sigma \in L} (L(\sigma) \sum_{|\sigma''| \leq K} \text{freq}(\sigma'_1 \cdot \sigma'' \cdot \sigma'_2, \sigma))} & \text{if } \exists \sigma \in L (\text{cov}(\sigma', \sigma'_1, \sigma'_2) \neq \emptyset) \\ 0 & \text{otherwise} \end{cases}$$

The numerator computes how many times σ' is covered by the context (σ'_1, σ'_2) in whole event log. The denominator computes how many times context (σ'_1, σ'_2) covers different substrings with length lower or equal than K . Clearly, the resulting value is a real number in $[0, 1]$. A higher value implies that it is more probable that subsequence σ' occurs among context (σ'_1, σ'_2) . So, $CP(\sigma', \sigma'_1, \sigma'_2, 1, L) = 1$, indicates that if context (σ'_1, σ'_2) occurs, subsequence σ' always happens among it. According to the event log that is presented in Table 1, $CP(\epsilon, \langle b \rangle, \langle c \rangle, 1, L) = \frac{7}{12}$ and $CP(\langle b \rangle, \langle b \rangle, \langle c \rangle, 1, L) = \frac{1}{12}$.

The main idea of our approach is that if the covering probability of a subsequence among a *significant context* is lower than the given threshold, the covered subsequence is considered as outlier. A significant context is simply a context that occurs significantly often, where the significance of occurrence is a user-defined threshold. Subsequently, we substitute the outlier subsequence with another subsequence that has a higher covering probability among the given context.

In context of our example, in the trace $\langle a, b, b, c, d, e, f, h \rangle$, if we consider $(\langle b \rangle, \langle c \rangle)$ as a frequent context, by replacing subsequence $\langle b \rangle$ with ϵ , which is more probable to happen among this context, we obtain $\langle a, b, c, d, e, f, h \rangle$, which is not an outlier any more. As another example, consider trace $\sigma = \langle a, d, b, d, c, f, h \rangle$. By having a significant context $(\langle a \rangle, \langle b \rangle)$, according to the whole event log, we expect that ϵ occurs among it rather than $\langle d \rangle$. Therefore, by this replacement we will have $\langle a, b, d, c, f, h \rangle$. Similarly, for $\langle a, b, c, d, g, h \rangle$ by considering context (d, g) and considering covering probabilities, we will replace ϵ by $\langle e \rangle$ and we obtain $\langle a, b, c, d, e, g, h \rangle$ which is not outlier any more.

The user decides which contexts are significant by setting a corresponding significance threshold. Contexts with a frequency of at least the number of traces in the event log times the given threshold are considered as significant contexts.

Also, the maximum length of covered subsequence (K) and context's subsequences (C_L) respectively are specified by the user. Note that C_L describes two values, i.e. a maximal length for σ'_1 and σ'_2 .

The input of our proposed method is an event log L , a context frequency threshold T_C , a context's subsequence lengths $C_L(l, r)$ and an upper bound for length of covered subsequence K . At first, all traces are scanned to compute covering probabilities of all contexts and their possible covered subsequences. Next, for each trace and each subsequence in it (with length $\leq K$), we check its context frequency and covering probability (according to T_C). If the context is significant and the covering probability is low, we replace the subsequence with a better one according to that context. Otherwise, if a context is insignificant it is not able to be used for repairing outlier behaviour.

A simple visual example of how the proposed method works is given in Fig. 2.

Significant contexts	Frequency	Probable subsequences
$(\emptyset, \langle a \rangle)$	98	\emptyset
$(\emptyset, \langle b \rangle)$	70	$\langle a \rangle$
$(\emptyset, \langle c \rangle)$	30	$\langle a \rangle$
$(\langle a \rangle, \langle b \rangle)$	101	$\emptyset, \langle c \rangle$
$(\langle a \rangle, \langle c \rangle)$	95	$\emptyset, \langle b \rangle$
$(\langle b \rangle, \langle c \rangle)$	67	\emptyset
$(\langle b \rangle, \emptyset)$	103	$\emptyset, \langle c \rangle$
$(\langle c \rangle, \langle b \rangle)$	35	\emptyset
$(\langle c \rangle, \emptyset)$	97	$\emptyset, \langle b \rangle$

$$L = \{ \langle a, b, c \rangle^{65}, \langle a, c, b \rangle^{30}, \langle a, b, b \rangle^3, \langle b, c \rangle^2 \}$$

$$\langle a, b, b \rangle \xrightarrow{\text{context}} \langle a, c, b \rangle$$

$$\epsilon \langle b, c \rangle \xrightarrow{\text{context}} \langle a, b, c \rangle$$

$$L = \{ \langle a, b, c \rangle^{67}, \langle a, c, b \rangle^{33} \}$$

Fig. 2. A simple example of repairing an event log L with our proposed method. Significant contexts and their probable subsequences are shown in the table for $K = 1$.

In this example, we consider $K = 1$ as a maximum subsequence length and also context lengths are equal to 1 to repair an event log with 100 traces. First, by scanning the event log, significant contexts and their probable subsequences are specified. If the corresponding context is significant and the subsequence is not probable for that context, we detect a noisy/infrequent behaviour. For example, the occurrence of $\langle b \rangle$ is improbable among significant context $(\langle a \rangle, \langle b \rangle)$.

After detecting outlier behaviour, we try to replace improbable subsequences with more probable ones. For substitution, we are searching for a subsequence with a length as close as possible to the outlier subsequence. Among all candidate subsequences we are interested in the subsequence with the highest probability. For example, if the outlier subsequence has length 2 we first search for a subsequence with the same length. If there is not any significant subsequence with length 2 for that context, we try to find a subsequence σ'' with length 1 or 3. Then, among the candidate subsequences we choose one that has the highest probability among that context. According to Fig. 2 there are two possible subsequences to substitute with $\langle b \rangle$, i.e. ϵ and $\langle c \rangle$. Because the length of $\langle c \rangle$ is similar

to the outlier subsequence we choose it and the trace is changed into $\langle a, c, b \rangle$. For the other outlier trace in this example, among the context $(\epsilon, \langle b \rangle)$, it is more probable that $\langle a \rangle$ occurs. So, by replacing ϵ with $\langle a \rangle$, it changes to $\langle a, b, c \rangle$.

To avoid having infinitive loops, after each replacement, the starting point of the next scanned subsequence will be the first point of the right subsequence of the previous context. For example, after replacement of $\langle a \rangle$ with ϵ and having $\langle a, b, c \rangle$, we will not check $\langle a \rangle$ again as a subsequence as we consider $\langle b \rangle$ as the next subsequence.

5 Evaluation

To evaluate our proposed method we conducted several experiments with both artificial and real event data. To simplify the evaluation, in all the experiments we just consider contexts' subsequences length equal to 1.

To apply the proposed repairing method, we implemented the *Repair Log* plug-in (*RL*) in the *ProM* framework¹. The plugin takes an event log as an input and outputs a repaired event log. Also, the user is able to specify threshold T_C , the maximum subsequence length K , and the length of left and right sequences of context C_L .

In addition, to apply our proposed method on various event logs with different thresholds and applying different process mining algorithms with various parameters, we ported the *Repair Log* plugin to *RapidProM*. *RapidProM* is an extension of *RapidMiner* that combines scientific workflows [22] with a range of (*ProM*-based) process mining algorithms.

To evaluate discovered process models, we use fitness and precision. Fitness computes how much behaviour in the event log is also described by the process model. However, precision measures how much of behaviour described by the model is also presented in the event log. Low precision means that the process model allows for much more behaviour compared to the event log. Note that, there is a trade off between these measures [23]. Sometimes, putting aside a small amount of behaviour causes a slight decrease in fitness, whereas precision increases much more. Therefore, to evaluate discovered process models, we use the F-Measures metric that combines fitness and precision: $\frac{2 \times \text{Precision} \times \text{Fitness}}{\text{Precision} + \text{Fitness}}$. Note that, in many applications, fitness has more importance. We therefore also use the notion of conditional precision, in which we only consider precision values of those process models that have $\text{fitness} \geq 0.95$. Furthermore, as shown in [4], *Matrix Filter* has a good performance on event logs with outlier behaviour. Hence, we compare the results of the proposed repairing method with this filtering method.

In the first experiment, we try to investigate the effect of using the proposed method on real event logs. Basic information of these event logs is presented in Table 3. We also add different percentages of noise to these event logs. As noise we consider addition of random activities at random positions in traces, random

¹ Repair Log plugin svn.win.tue.nl/repos/prom/Packages/LogFiltering.

Table 3. Details of real event logs that are used in the experiment

Event log	Activities count	Traces count	Events count	Variants counts
<i>BPIC_2017_Offer</i>	8	42995	193849	16
<i>BPIC_2013</i>	13	7554	65533	2278
<i>BPIC_2012_Application</i>	10	13087	60849	17
<i>BPIC_2012_Workflow</i>	6	9658	72413	2263
<i>BPIC_2012_Offer</i>	7	5015	31244	168
<i>Road_Fines</i>	11	150370	561470	231
<i>Hospital_Billing</i>	18	100000	451359	1020
<i>Sepsis_Cases</i>	16	1050	15224	846
<i>Credit</i>	15	10035	150525	1

removal of activities, and swapping of activities within traces. For example, we added 5% of all these three mentioned types of noise to *Road_Fines* and refer to it as *Road_Fines_05*. Note that, in all experiments, we just used the modified (filtered or repaired) event logs for discovering process models. For conformance checking and quality assessment of discovered process models we used original event logs (without added noise).

We try to discover the best process models according to the F-Measure for these event logs using four different methods. In the first method, we apply *IMi* [11] with 51 different noise filtering thresholds (that is embedded in the Inductive Miner) from 0 to 1 and we call it *N&IMi* (*N* means that event logs are not filtered beforehand). In the second method, we apply the basic Inductive Miner (without its embedded filtering) on an event log that is previously filtered with *Matrix Filter* method and call it as *M&IM* (*M* means event logs are filtered with *Matrix Filter* beforehand). We also apply the basic Inductive Miner on event logs that are repaired beforehand and name it as *R&IM* (*R* means event logs are repaired with the proposed method beforehand). In the last method, the Inductive Miner with four different noise filtering thresholds (from 0.1 to 0.4) is applied on repaired event logs (*R&IMi*).

The results of applying these methods to aforementioned event logs and their best F-Measure values are given in Figs. 3 and 4. These figures show that when there is outlier behaviour in the event log, just using the Inductive Miner (*N&IMi*) does not yield a proper process model. However, in some event logs that do not contain outlier behaviour, the Inductive Miner discovers a process model with a suitable F-Measure. It is also notable that for most event logs, *R&IM* outperforms *M&IM* and allows us to obtain better results. But, for the *Hospital_Billing* event log, the result of *M&IM* is slightly better in the previous experiments. This is because there are lots of variants in this event log, but 94% of traces are placed in just 1% of variants and in this type of event logs, filtering works perfect. The best results are achieved when we apply Inductive

Miner on cleaned event logs or *R&IMi*. Moreover, in most of cases, *M&IM* and the Inductive Miner (here *N&IMi*) are achieving their best possible process model according to F-Measure, by sacrificing a lot in term of fitness. In Figs. 5 and 6, we see the result of different methods according to their conditional precision. According to these figures, the *R&IM* method helps the Inductive Miner to discover process models with high precision without sacrificing the fitness. However, we see that for event logs with lots of outlier behaviour the Inductive Miner does not able to discover a process model with high fitness and precision at the same time.

Note that for some event logs like the *Sepsis_Cases* neither *M&IM* nor *R&IM* allow us to discover an outstanding process model, i.e., having high conditional precision. The reason is for this is related to the fact that the discovered process model of this event log is not precise and there is some behaviour that possible to happen anywhere during execution of the process. We have shown the best discovered process model for this event log to a business expert which indicated that the process model is acceptable (but not perfect).

In Fig. 7, the best process model of each methods for the *BPIC_2017_Offer_05* are shown. It is obvious that Fig. 7a that is discovered on the repaired event log is the best process model among them.

Note that, in the previous experiments, the original event log that is used to compute F-Measures (our ground truth) potentially also contains noisy and infrequent behaviour. Therefore, in another experiment, we create six artificial process models with different behaviour. These process models consequently describe different types of behaviour, i.e. one model just contains sequence constructs (*Sequence*), one contains sequence and choice constructs (*Xor*), one con-




Fig. 3. Best F-measure of applying different methods on *BPIC* event logs.




Fig. 4. Best F-measure of applying different methods on other real event logs.




Fig. 5. Best conditional precision of applying different methods on *BPIC* event logs.






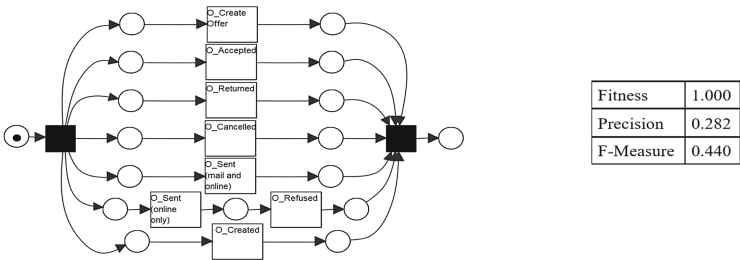
Fig. 6. Best conditional precision of applying different methods on other real event logs.



(a) Results of Applying *R&IM*



(b) Results of applying *M&IM*



(c) Results of applying *N&MI*

Fig. 7. The best discovered process models on *BPIC-2017-Offer_05* using *RL*, *MF* and *NF* methods.




Fig. 8. Best F-measure of applying different methods on synthetic event logs.




Fig. 9. Best conditional precision of applying different methods on synthetic event logs.

tains sequence and parallel constructs (*Parallel*), one contains sequence and loop constructs (*Loop*), one contains all possible behavioural constructs (*All*) and we use a process model with locally structured subprocesses that has lots of variants (*High Variants*). Based on these reference process models, we generated six event logs that each contains 5000 traces. Similar to the previous experiment, 5%, 10%, 20% and 50% of noise inserted to these original event logs. Note that, an event log with 10% inserted noise not necessary has 90% original traces. Similarly, we apply *N&IMi*, *M&IM*, *R&IM* and *R&IMi* to these event logs. Results of this experiment are given in Figs. 8 and 9.

As shown in these figures, for all event logs if we insert some noisy behaviour (even just 5%), the Inductive Miner has problems to discover a process model with a high F-Measure. This means the embedded noise filtering mechanism in this algorithm is not able to deal with all types of outlier. Furthermore, *Matrix Filter* always works better than the embedded noise filtering method in the Inductive Miner. Compared to these filtering methods, the proposed repaired method performs better especially when there probability of noise is low. Only for the *Parallel* event log, results of applying *M&IM* slightly outperforms our proposed method. Like the previous experiment, the best results are achieved when we apply the Inductive Miner with its filtering method to the repaired event logs (*R&IMi*). Except for the *Parallel* event logs, for all others this method finds the underlying process models, i.e. almost equal to the reference process models. For the *High Variants* event log that has local structured subprocesses, our proposed repair method works perfectly and accurately detects and repairs outlier behaviour. The reason for this is related to the fact that our method is able to detect and repair outlier behaviour locally.

Table 4. The percentage of remaining traces in event logs for the best process model

Event Log Noise Percentage	0	5	10	20	50
Sequence	100	95	90	81	52
Xor	100	95	90	42	52
Parallel	100	52	48	41	23
Loop	100	95	90	80	51
ALL	83	95	90	9	5
High Variants	100	95	90	81	20

Finally, note that filtering methods achieve the best results by removing a lot of behaviour in event logs. The percentages of remaining traces in each event log for the best model in *M&IM* are given in Table 4. In some cases the best process model is discovered just with 5% of the traces. This means that we need to remove a lot of behaviour from the event log. However, in the repair method all the traces remain in the event log but they may be modified.

Note, for all methods we use a grid search on different parameters and show the best results obtained. However, in reality like other state-of-the-art process mining specific data cleansing methods, adjusting these thresholds is a challenging task for users.

6 Conclusion

Process mining provides insights into the actual execution of business processes, by exploiting available event data. Unfortunately, most process mining algorithms are designed to work under the assumption that the input data is outlier free. However, real event logs contain outlier (noise and infrequent) behaviour that typically leads to inaccurate/unusable process mining results. Detecting such behaviour in event logs and correcting it helps to improve process mining results, e.g. discovered process models.

To address this problem, we propose a method that takes an event log and returns a repaired event log. It uses the occurrence frequency of a control-flow oriented context pattern and the probabilities of different subsequences appearing in middle of it to detect outlier behaviour. If such probability be lower than a given threshold, the subsequence is substituted with a more probable one according to the context.

To evaluate the proposed repairing method we developed a plug-in in the ProM platform and also offer it through RapidProM. As presented, we have applied this method on several real event logs, and compared it with other state-of-the-art process mining specific data cleansing methods. Additionally, we applied the proposed method on synthetic event logs. The results indicate that the proposed repair approach is able to detect and modify outlier behaviour and consequently is able to help process discovery algorithms to return models

that better balance between different behavioural quality measures. Furthermore, using these experiments we show that our repair method outperforms one of the best state-of-the-art process mining filtering techniques as well as the Inductive Miner algorithm with its embedded filtering mechanism.

As future work, we want to evaluate the proposed method on other process mining results and also other process discovery algorithms. We also plan to develop techniques to automatically set adjustable filtering and repairing parameters based on characteristics of the input event log to guide users and speed-up analysis.

References

1. van der Aalst, W.M.P.: Using process mining to bridge the gap between BI and BPM. *IEEE Comput.* **44**(12), 77–80 (2011)
2. van der Aalst, W.M.P.: *Process Mining - Data Science in Action*, 2nd edn. Springer, Heidelberg (2016)
3. Conforti, R., La Rosa, M., ter Hofstede, A.H.M.: Filtering out infrequent behavior from business process event logs. *IEEE Trans. Knowl. Data Eng.* **29**(2), 300–314 (2017)
4. Sani, M.F., van Zelst, S.J., van der Aalst, W.M.P.: Improving process discovery results by filtering outliers using conditional behavioural probabilities. In: Teniente, E., Weidlich, M. (eds.) *BPM 2017. LNBP*, vol. 308, pp. 216–229. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-74030-0_16
5. van der Aalst, W., van Dongen, B.F., Günther, C.W., Rozinat, A., Verbeek, E., Weijters, T.: ProM: the process mining toolkit. *BPM (Demos)* **489**(31) (2009)
6. van der Aalst, W.M.P., Bolt, A., van Zelst, S.J.: RapidProM: mine your processes and not just your data. *CoRR* abs/1703.03740 (2017)
7. van der Aalst, W., et al.: Process mining manifesto. In: Daniel, F., Barkaoui, K., Dustdar, S. (eds.) *BPM 2011. LNBP*, vol. 99, pp. 169–194. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-28108-2_19
8. Rebuge, Á., Ferreira, D.R.: Business process analysis in healthcare environments: a methodology based on process mining. *Inf. Syst.* **37**(2), 99–116 (2012)
9. van der Aalst, W.M.P., Weijters, T., Maruster, L.: Workflow mining: discovering process models from event logs. *IEEE Trans. Knowl. Data Eng.* **16**(9), 1128–1142 (2004)
10. Leemans, S.J.J., Fahland, D., van der Aalst, W.M.P.: Discovering block-structured process models from event logs - a constructive approach. In: Colom, J.-M., Desel, J. (eds.) *PETRI NETS 2013. LNCS*, vol. 7927, pp. 311–329. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-38697-8_17
11. Leemans, S.J.J., Fahland, D., van der Aalst, W.M.P.: Discovering block-structured process models from event logs containing infrequent behaviour. In: Lohmann, N., Song, M., Wohed, P. (eds.) *BPM 2013. LNBP*, vol. 171, pp. 66–78. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-06257-0_6
12. van Zelst, S.J., van Dongen, B.F., van der Aalst, W.M.P., Verbeek, H.M.W.: Discovering workflow nets using integer linear programming. *Computing* (2017)
13. Weijters, A.J.M.M., Ribeiro, J.T.S.: Flexible heuristics miner (FHM). In: *CIDM* (2011)

14. Günther, C.W., van der Aalst, W.M.P.: Fuzzy mining – adaptive process simplification based on multi-perspective metrics. In: Alonso, G., Dadam, P., Rosemann, M. (eds.) BPM 2007. LNCS, vol. 4714, pp. 328–343. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-75183-0_24
15. Chandola, V., Banerjee, A., Kumar, V.: Anomaly detection for discrete sequences: a survey. *IEEE Trans. Knowl. Data Eng.* **24**(5), 823–839 (2012)
16. Wang, J., Song, S., Lin, X., Zhu, X., Pei, J.: Cleaning structured event logs: a graph repair approach. In: ICDE 2015, pp. 30–41 (2015)
17. Cheng, H.J., Kumar, A.: Process mining on noisy logs-can log sanitization help to improve performance? *Decis. Support Syst.* **79**, 138–149 (2015)
18. van Zelst, S.J., Fani Sani, M., Ostovar, A., Conforti, R., La Rosa, M.: Filtering spurious events from event streams of business processes. In: Proceedings of the CAISE (2018)
19. Fahland, D., van der Aalst, W.: Model repair-aligning process models to reality. *Inf. Syst.* **47**, 220–243 (2015)
20. Armas-Cervantes, A., van Beest, N., La Rosa, M., Dumas, M., Raboczi, S.: Incremental and interactive business process model repair in Apromore. In: Proceedings of the BPM Demos. CRC Press (2017)
21. Rogge-Solti, A., Mans, R.S., van der Aalst, W.M.P., Weske, M.: Improving documentation by repairing event logs. In: Grabis, J., Kirikova, M., Zdravkovic, J., Stirna, J. (eds.) PoEM 2013. LNBIP, vol. 165, pp. 129–144. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-41641-5_10
22. Bolt, A., de Leoni, M., van der Aalst, W.M.P.: Scientific workflows for process mining: building blocks, scenarios, and implementation. *STTT* **18**(6), 607–628 (2016)
23. Weerd, J.D., Backer, M.D., Vanthienen, J., Baesens, B.: A robust f-measure for evaluating discovered process models. In: Proceedings of the CIDM, pp. 148–155 (2011)

ICT Project Management



Big Data Enabled Organizational Transformation: The Effect of Inertia in Adoption and Diffusion

Patrick Mikalef¹✉, Rogier van de Wetering², and John Krogstie¹

¹ Department of Computer Science, Norwegian University of Science and Technology, Sem Saelandsvei 9, 7491 Trondheim, Norway
{patrick.mikalef, john.krogstie}@ntnu.no

² Faculty of Management Science and Technology,
Open University of the Netherlands, 6401 DL Heerlen, The Netherlands
rogier.vandewetering@ou.nl

Abstract. Big data and analytics have been credited with being a revolution that will radically transform the way firms operate and conduct business. Nevertheless, the process of adopting and diffusing big data analytics, as well as actions taken in response to generated insight, necessitate organizational transformation. Nevertheless, as with any form of organizational transformation, there are multiple inhibiting factors that threaten successful change. The purpose of this study is to examine the inertial forces that can hamper the value of big data analytics throughout this process. We draw on a multiple case study approach of 27 firms to examine this question. Our findings suggest that inertia is present in different forms, including economic, political, socio-cognitive, negative psychology, and socio-technical. The ways in which firms attempt to mitigate these forces of inertia is elaborated on, and best practices are presented. We conclude the paper by discussing the implications that these findings have for both research and practice.

Keywords: Big data analytics · Organizational transformation
Inertia · Deployment · IT-enabled transformation

1 Introduction

While big data analytics have been in the spotlight of attention by researchers and practitioners in the last few years, to date there has been limited attention on what forces can potentially hinder the potential business value that these investments can deliver. Much research has focused on the necessary investments that must be made to derive business value [1], but the process from making the decision to adopt such technologies, up to turning insight into action is seldom discussed, particularly with respect to inertia. The underlying premise of big data dictates that such investments can generate insight with the potential to transform the strategic direction of firms, and help them outperform competition [2]. Nevertheless, this process entails organizational transformation at multiple levels, and as with any case of organizational transformation, is subject to path dependencies, routinization, and other hindering forces [3].

While big data literature has documented the importance that organizational learning and a data-driven culture have on overall project success [4], there is to date a very limited understanding on how these should be implemented and what factors may inhibit successful deployment or even adoption. In this respect, there is not much attention on the processes of big data adoption and implementation. Most studies to date have attempted to provide a narrative on how big data can produce value [5], or even empirically show an association between investments and performance measures [1, 6]. Yet, in reality, managers and practitioners are faced with a number of hurdles which need to be overcome, on individual, group, organizational, and industry levels. The purpose of this study is therefore to attempt to understand how inertial forces in these levels hinder the potential value of big data analytics. By doing so, it is possible to isolate key success factors of implementation and help guide practitioners in developing strategies for adoption and deployment [7].

Hence, this research is driven by the following research question which helps guide our investigation:

How is inertia present in big data projects? At what stages do inertial forces appear and at what levels?

To answer these questions, we build on the extant literature on organizational transformation and on studies focusing on inertia in IT-based implementations. We isolate five key forms of inertia, namely, *economic, political, socio-cognitive, negative psychology, and socio-technical* and examine how each of these is present in big data analytics projects in firms. Following a multiple case study approach in which we interview higher level executives of IT departments from 27 firms, we present findings and discuss the implications that they create for both research and practice.

2 Big Data Analytics and Business Value

Big data analytics is widely regarded as the next frontier of improving business performance due to its high operational and strategic potential [8]. The literature defined big data analytics as “a new generation of technologies and architectures, designed to economically extract value from very large volumes of a wide variety of data, by enabling high velocity capture, discovery and/or analysis” [9]. While most claims on the value of big data analytics are anecdotal, the emerging literature has documented a positive relationship between the decision to invest in firm-wide deployment of big data analytics and performance [1]. Big data analytics enable firms to make sense of vast amounts of data and reconfigure their strategies based on trends that are observed in their competitive environment [10]. The importance of big data analytics is evident from the increasing investments made from firms, and particularly those working in complex and fast-paced environments [11]. Managers nowadays are relying ever more on big data analytics to inform their decision-making and direct future strategic initiatives [12]. The value of investing in big data analytics is clearly reflected in a recent article by Liu [13], who notes that big data analytics constitutes a major differentiator between high-performing and low-performing firms, as it enables firms to be more proactive and swift in identifying new business opportunities. Additionally, the study

reports that big data analytics have the potential to decrease customer acquisition costs by 47% and enhance revenues by about 8%. A report by MIT Sloan Management Review shows that companies that are leaders in the adoption of big data analytics are much more likely to produce new products and services compared to those that are laggards [14]. Nevertheless, the value that firms realize from big data investments, is highly contingent upon the idiosyncratic capabilities that they develop in deriving meaningful insight [15, 16].

3 Organizational Inertia and Decision Making

Understanding what factors facilitate or inhibit organizational adoption and diffusion of emerging information technologies (IT) has long been a primary concern for researchers and practitioners [17]. The main premise associated with adoption of any new IT innovation is that it entails a level of organizational transformation to both incorporate IT into operations as well as improve business performance as a result of it [18]. Nevertheless, it is frequently noted that when transformation is required, organizations are rigid and inert, resulting in the overall failure of the newly adopted IT [19]. Past studies in management science and in the information systems literature have examined and isolated different forms of inertia manifested at different levels and throughout numerous agents [20]. Nevertheless, despite several studies examining the effect of inertia on different contexts and types of IT, there is very limited research on the role big data analytics play, and the inertial forces that can possibly slow down implementation and hinder business value. To understand these and derive theoretical and practical implications, we first start by surveying the status quo of existing literature on organizational inertia, particularly with regards to IT adoption and diffusion.

Organizational inertia is a subject that has long been in the center of attention for scholars in the managerial science domain. Inertia represents the price for stable and reproducible structures that guarantee the derided reliability and accountability of organizations [21]. Nevertheless, the presence of inertia is usually discernible in the need for change and is distinguishable when external stimuli demands so. The process of realigning the organization with the environment therefore requires that the forces of inertia that are present within an organization need to be overcome [18]. We build on the extant literature in the domain of IT-enabled organizational transformation and management science and identify five broad types of inertia [18, 22–24]. These include negative psychology inertia, socio-cognitive inertia, socio-technical inertia, economic inertia, and political inertia [18]. In the context of IT research, Besson and Rowe [18] provide a clear definition of what inertia represents in the face of novel organizational implementation. Specifically, they state that “*inertia is the first level of analysis of organizational transformation in that it characterizes the degree of stickiness of the organization being transformed and defines the effort required to propel IS enabled organizational transformation*”. They do however mention that identifying the sources of inertia is only one level, the second being process and agency, and the third performance. These levels help distinguish causes of inertia from strategies to overcome them and quantifiable measures to assess their impact on organizational transformation.

The first step however is to clearly define and understand how the different sources of inertia have been examined in literature and at what level they are present. Negative psychology inertia has been predominantly attributed to group and individual behavior, and is based on threat perceptions of losing power or even their position. Uncertainty about the role of individuals or groups in the face of novel technological deployments thus causes negative psychological reactions which biases them towards the *status quo* [25]. Socio-cognitive inertia emphasize mostly on malleability due to path dependencies, habitualization, cognitive inertia and high complexity [26]. These forms of inertia arise due to periods of sustained stability and routinization caused by a stable environment. Socio-technical inertia on the other hand refers to the dependence on socio-technical capabilities, which arise from the interaction of the social systems and technical system and their joint optimization [27]. Economic inertia may be present in the form of commitment to previously implemented IT solutions that do not pay off and create sunk costs, or through transition expenses which cause organizations to not adopt potentially better alternatives [19]. Finally, political inertia is caused by vested interests and alliances which may favor that the organization remains committed to a specific type of information technology so that partnerships are not broken.

While to date there has been no systematic study to examine the forms of inertia in big data analytics implementations, several research studies have reported inhibiting factors during adoption and diffusion. Mikalef, Framnes, Danielsen, Krogstie and Olsen [28] mention that in some cases economic inertia caused a problem in the adoption of big data analytics. The authors state that top managers in some cases were reluctant to make investments in big data analytics, since their perceptions about the cost of such investments in both technical and human resources greatly exceeded the potential value. In addition, they mention that both socio-cognitive and socio-technical issues rose at the group level, where people were reluctant to change their patterns of work, and were also afraid of losing their jobs. Similar findings are reported by Janssen, van der Voort and Wahyudi [15], where socio-cognitive inertia can be reduced by implementing governance schemes [29], which dictate new forms of communication and knowledge exchange. In their study, Vidgen, Shaw and Grant [4] note that inertial forces impact the implementation of big data projects, and that the presence of the right people that can form data analytics teams and implement processes is critical to success.

4 Research Method

4.1 Design

Beginning from the theoretical background and the overview of existing literature on big data-enabled organizational transformation and business value, the present work aims to understand how the processes of deploying big data analytics within companies is hindered by different forms of inertia as well as decision making barriers by top managers in the process of decision making. We explain how inertia is presented at different forms and stages throughout the deployment and routinization of big data

analytics projects. Specifically, we base our investigation on the following research question:

What hindrances are detectable during the process of big data-driven organizational transformation? At which stages are they detectable and how can they be overcome?

We started our investigation by surveying past literature on the main challenges associated with IT-enabled organizational transformation. The purpose of this review was to understand the primary reasons IT solutions fail to deliver business value. Next, we attempted to understand how these notions are relevant to companies that have initiated deployments of big data analytics projects. To do this, this study followed a multiple case-study approach. The case study methodology is particularly well-suited for investigating organizational issues [30]. Through multiple case studies, we are able to gain a better understanding of the frictions that are created between different employees and business units during the implementation of big data analytics, as well as the causes of non-use of generated insight by top managers. Through a multiple case study approach, it is also possible to enable a replication logic in which the cases are treated as a series of experiments that confirm or negate emerging conceptual insights [31]. We chose a deductive multiple case study analysis based primarily on interviews with key informants, and secondary on other company-related documents. This selection was grounded on the need to sensitize concepts, and uncover other dimensions that were not so significant in IT-enabled organizational transformation studies [32].

4.2 Research Setting

For the sample of companies that are included in our multiple case study approach, we selected firms that demonstrated somewhat experience with big data analytics. This included companies that had either just recently started or had invested considerable time and effort in gaining value from big data. In addition, we focused mostly on medium to large size companies since the complexity of the projects they were involved in would give us a better understanding of the spectrum of requirements in big data initiatives. Lastly, the firms we selected operated in competitive and highly dynamic markets which necessitated the adoption of big data as a means to remain competitive. These companies also faced mimetic pressures to adopt big data since in most cases they were afraid that competitors would overtake them if they did not follow the big data paradigm. Therefore, efforts in developing strong organizational capabilities via means of big data analytics were accelerated. We selected different companies in terms of type of industry within the given boundaries, with the aim of doing an in-depth analysis and to be in place to compare and contrast possible differences (Table 1). The selected firms are considered established in their market in the European region, with most companies being based in Norway, the Netherlands, Italy, and Germany.

Table 1. Profile of firms and respondents

Code	Business areas	Employees	Primary objective of adoption	Key respondent (Years in firm)
C.1	Consulting services	15.000	Risk management	Big Data and Analytics Strategist (4)
C.2	Oil & Gas	16.000	Operational efficiency, Decision making	Chief Information Officer (6)
C.3	Media	7.700	Market intelligence	Chief Information Officer (3)
C.4	Media	380	Market intelligence	IT Manager (5)
C.5	Media	170	Market intelligence	Head of Big Data (4)
C.6	Consulting services	5.500	New service development, Decision making	Chief Information Officer (7)
C.7	Oil & Gas	9.600	Process optimization	Head of Big Data (9)
C.8	Oil & Gas	130	Exploration	IT Manager (6)
C.9	Basic Materials	450	Decision making	Chief Information Officer (12)
C.10	Telecommunications	1.650	Market intelligence, New service development	Chief Digital Officer (5)
C.11	Financials	470	Audit	IT Manager (7)
C.12	Retail	220	Marketing, Customer intelligence	Chief Information Officer (15)
C.13	Industrials	35	Operational efficiency	IT Manager (5)
C.14	Telecommunications	2.500	Operational efficiency	IT Manager (9)
C.15	Retail	80	Supply chain management, inventory management	Chief Information Officer (11)
C.16	Oil & Gas	3.100	Maintenance, Safety	IT Manager (4)
C.17	Technology	40	Quality assurance	Head of IT (3)
C.18	Technology	180	Customer management, Problem detection	IT Manager (7)
C.19	Oil & Gas	750	Decision making	Chief Information Officer (14)
C.20	Technology	8	Business intelligence	Chief Information Officer (3)
C.21	Basic materials	35	Supply chain management	Chief Information Officer (6)
C.22	Technology	3.500	New business model development	Chief Digital Officer (8)
C.23	Technology	380	Personalized marketing	IT Manager (2)
C.24	Basic materials	120	Production optimization	IT Manager (4)
C.25	Technology	12.000	Customer satisfaction	Chief Information Officer (15)
C.26	Technology	9	Product function, machine learning	Chief Information Officer (2)
C.27	Telecommunications	1.550	Fault detection, Energy preservation	Chief Information Officer (9)

4.3 Data Collection

In this study, we collected data from primary sources, as well as secondary sources to confirm statements and establish robustness. The primary sources were the direct interviews that were conducted with key respondents in firms. The interview focused on their attitudes, beliefs, and opinions regarding their experience with big data initiatives that their firm had undertaken. All interviews were conducted face-to-face in a conversational style, opening with a discussion on the nature of the business and then proceeding on to the themes of the interview guideline. Overall a semi-structured case study protocol was followed in investigating cases and collecting data [33]. Discussion were recorded and then transcribed for analysis. Two of the co-authors completed the independent coding of the transcripts in accordance with the defined themes as identified in Table 2. Each coder read the transcripts independently to find specific factors related to the types of inertia, as well as on biases of managers in making insight-driven decisions and the reasons they do so. This process was repeated until inter-rater reliability of the two coders was greater than 90% [34].

4.4 Data Analysis

The empirical analysis was performed by an iterative process of reading, coding, and interpreting the transcribed interviews and observation notes of the 27 case studies [35]. At a first stage we identified and isolated a large number of concepts based on the literature that was discussed in earlier sections. For each case the standardization method was used to quantify these characteristics using an open coding scheme [33]. This allowed us to cluster primary data in a tabular structure, and through the iterative process identify the relative concepts and notions that were applicable for each case. Collectively, these concepts (Table 2) comprise what is referred to in literature as organizational inertia [18]. The underlying logic suggests that there are multiple barriers when examining the value of big data projects of firm performance, some of which are caused due to organizational inertia and are discernible at different stages of implementation, while others appear at the decision-making stage, in which managers for a combination of reasons tend not to adopt the insight that is generated by big data analytics, but rather follow their instinct [9]. The effect of a firms' big data analytics capability on performance is therefore considered to be mediated and moderated by numerous factors that appear at different stages of the implementation process.

5 Results

After transcribing the interviews and assigning them each a thematic tag as those described in Table 2, we started aggregating finding and identifying common patterns. More specifically, the inertial forces and how they are presented in big data projects are summarized below.

Table 2. Thematic support for organizational inertia

Inertia dimensions	Perspective of agent	Level	References
Economic	Agents are embedded in business models that have their own dynamics arising from resource reallocation between exploitation and exploration processes	Business and sector	Besson and Rowe [18], Kim and Kankanhalli [25]
Political	Agents are embedded in networks of vested interests that have their own dynamics, especially due to alliances rebuilding time	Business	Besson and Rowe [18]
Socio-cognitive	Agents are embedded in institutions characterized by their stickiness due to norms and values re-enactment	Individual, group, organization and industry	Besson and Rowe [18], Haag [19]
Negative psychology	Agents are overwhelmed by their negative emotions due to threat perception	Individual and group	Besson and Rowe [18], Polites and Karahanna [20]
Socio-technical	Agents are embedded in socio-technical systems that have their own dynamics, especially due to development time and internal consistency	Group and organization	Besson and Rowe [18], Lyytinen and Newman [26]

5.1 Economic Inertia

Economic inertia was a very prominent theme amongst most of the companies, especially those that were not multi-national firms or had major slack resources, such as micro, small, and medium enterprises. For large conglomerate companies, scarcity of economic resources towards the implementation of big data projects was not an issue. Specifically, in non-large companies', economic inertia was present from the top management and board of directors, who had doubts about the value of big data analytics in their operations. The respondents from companies C.13 and C.24 respectively made the following comments.

“For us the value of big data analytics was not clear. We did not want to invest in a fashion just because everyone else is doing it. These technologies are expensive because we need personnel that we currently do not have.”

“The management was very skeptical about if we should go into big data. Our competitors were doing that and it meant that we had to follow so that we are not left behind. We tried to experiment in the beginning with some internal resources but then we understood that we have to invest more. This was a hard decision to make and perhaps we delayed a bit on this (adopting big data)”

Similar quotes were made by several other executives, showcasing that economic inertia are a major inhibiting factor of big data adoption and deployment. A major

inhibitor that leads to this is the unclear link between big data investments and business value. On the other hand, competitive pressures seem to be driving mimetic behaviors in companies mitigating the effects of economic inertia.

5.2 Political Inertia

Political inertia was detected in several firms that had formed partnerships with other private and public organizations. Specifically, in C.5 the manager made some remarks about lock-in effects that vendors of information systems led them to. The IT department of C.5 wanted to utilize data of their own in combination with some from a partner of a hospital. The vendor of the partner hospital information systems didn't allow for the extraction and use of data in third party analytics tools, and promoted his company's analytics tools. The top management of the hospital wanted to retain good relationships with the vendor of their information systems because they formed the backbone of operations for different departments, so the partnership between the two entities was forced to collapse. Similar phenomena were observed especially in the case of private-public alliances, where public bodies are trapped by vested interests.

5.3 Socio-Cognitive Inertia

Socio-cognitive inertia were found to be a problem in most of the companies that were examined. In most of the cases, big data implementation meant that data from different departments needed to be gathered. This entailed that a detailed account of what data were available should be initiated, and in many cases new processes needed to be put into place to collect data. Typically, the IT department was responsible for starting this process, and had to explain to various other siloed departments their goals, ways of realizing them, and what their role would be. The different mental modes, use of language, and objectives caused conflicts that threatened and even greatly delayed big data implementation projects. An example of this is evident through the comments of respondent of C.10.

"I think we had some major issues when we went into the marketing department and asked them to talk to us about how they gather data on customer preferences based on advertisement campaigns. We understood that they didn't have any feedback mechanism in place to evaluate the success or failure of what they did. This meant that somehow, they had to track the feelings and attitudes of consumers. I think that we also confused them about the data we were looking to collect, this confusion also led to a bit of tension"

Similar examples were found in several companies working external partners such as universities, public bodies, and other companies, as well as in firms that have highly siloed business units. We found that in many cases, consulting firms were brought in to resolve this issue and act as a mediating agent. They tried to create a common understanding of the objectives of the big data projects, and bring representatives from each business unit to the table so that cognitive structures can be in alignment.

5.4 Negative Psychology

Negative psychology was again observed mostly in small and medium firms, where the IT department comprised of a small number of employees. Primarily, it was found in personnel that had been actively employed for many years, in contrast to those that had recently graduated. These employees feared that the introduction of big data analytics and the corresponding technologies and tools for analyzing and visualizing data would render their skills as non-significant. Specifically, the respondent from C.8 stated the following.

“When I told the group that we should start to think about what we can do with the data in the company there was a cold silence. They initially objected saying that there would be not much value in doing so and that it was just a waste of time. Others said that they didn’t have enough time to engage in this process since it would take a lot of time to learn how these technologies worked. I realized that deep inside they feared that the move to big data would require for them to learn new stuff, or even mean that they could lose their jobs if they didn’t manage to adapt”

We saw that the way many IT managers handled this issue was by providing their division with small challenges, and incrementally growing projects. Also, they assigned a few hours a week where they had the freedom to experiment with big data tools. This allowed them to try out these novel tools at their own pace, without the fear of time-pressure to deliver results.

5.5 Socio-Technical

In terms of socio-technical inertia, it was observed that in many cases middle-level managers exerted behaviors that stalled the implementation of big data projects. Their primary fear was that decision-making would now reside in insight from analytics, therefore replacing them. In many occasions this fear manifested itself as a distrust towards the value of big data analytics, and general tendency to downplay the significance of big data in operations. Despite the clear directive of top management to diffuse big data analytics into operations, in many circumstances middle-level managers would not invest time in clear implementation strategies leaving the IT department with a bleak understanding of how they should proceed. The respondent from C.6 states the following.

“When we made the decision to start working with big data we were happy. Then after some time we were not sure what we should do with it. We presented some examples to management but they didn’t take them much into account. We then realized that they prefer to make decisions based on their own knowledge, and didn’t trust the process we were following, or even, felt like it may replace them. It was quite discouraging to work on something that is not applied”

Typically, these issues were resolved by a strong top management vision and leadership. In addition, training seminars for middle-level managers on the value of big data and their role were regarded as very beneficial in overall success.

6 Discussion

In the current study we have examined how inertia during the implementation and deployment phases of big data projects influence their success. Following the literature which distinguishes between economic, political, socio-cognitive, negative psychology, and socio-technical, we looked at how these forces of inertia are manifested in contemporary organizations through 27 case studies. Our results show that value from big data investments, and even actual implementation, can be hindered by multiple factors and at multiple levels which need to be considered during the planning phase. To the best of our knowledge this is one of the first attempts to isolate these inhibiting forces and provide suggestions on which future research can build and managers can develop strategies for adopting and diffusing their big data investments.

From a research perspective the major finding of this research is that even in the presence of all necessary big data analytics resources, there are multiple ways in which a business value can be hindered. This raises the question of how can these obstacles be overcome. While there is a stream of research into the issues of information governance these studies primarily focus on the issue of how to handle data and how to appropriate decision making authority in relation to the data itself. There still seems to be an absence of governance schemes that follow a holistic perspective and include management and organization of all resources, including human and intangible ones. In addition, how firms should handle individual, group and industry-level dynamics is a topic that is hardly touched upon.

From a managerial point of view, the outcomes of this study outline strategies that can be followed to mitigate the effects of the different types of inertia. Our findings indicate that inertia can be present at many phases of adoption and diffusion so action need to be taken throughout projects. It is critical to consider the socio-technical challenges that these technologies create for middle-level managers and clearly understand how their decision-making is influenced or not by insight generated by big data. In addition, it is important to develop strategies so that the whole organization adopts a data-driven logic, and that a common understanding and language is established. With regards to the IT department, educational seminars and incremental projects seem to be the way to limit negative psychology barriers. Also, providing a clear sense of direction as to what kind of analytics are to be performed on what data is of paramount importance. It is commonly observed that many companies delve into the hype of big data without having a clear vision of what they want to achieve.

While this research helps to uncover forces of inertia and the levels at which they present themselves, it does not come without limitations. First, we looked at companies that have actually adopt big data, a more complete approach would be to look at what conditions cause other firms to not opt for big data. Second, while we briefly touched on the issue of middle-level managers not following insight generated from big data, it is important to understand in more detail the decision-making processes that underlie their reasoning. Also, the actions that are taken in response to these insights are seldom put into question. This is a future area which should be examined since the value of big data cannot be clearly documented in the absence of knowledge about strategic or operational choices.

Acknowledgments.

This project has received funding from the European Union's Horizon 2020 research and innovation programme, under the Marie Skłodowska-Curie grant agreement No. 704110.

References

1. Gupta, M., George, J.F.: Toward the development of a big data analytics capability. *Inf. Manag.* **53**, 1049–1064 (2016)
2. Prescott, M.: Big data and competitive advantage at Nielsen. *Manag. Decis.* **52**, 573–601 (2014)
3. Sydow, J., Schreyögg, G., Koch, J.: Organizational path dependence: opening the black box. *Acad. Manag. Rev.* **34**, 689–709 (2009)
4. Vidgen, R., Shaw, S., Grant, D.B.: Management challenges in creating value from business analytics. *Eur. J. Oper. Res.* **261**, 626–639 (2017)
5. McAfee, A., Brynjolfsson, E., Davenport, T.H.: Big data: the management revolution. *Harvard Bus. Rev.* **90**, 60–68 (2012)
6. Wamba, S.F., Gunasekaran, A., Akter, S., Ren, S.J.-f., Dubey, R., Childe, S.J.: Big data analytics and firm performance: effects of dynamic capabilities. *J. Bus. Res.* **70**, 356–365 (2017)
7. Mikalef, P., Pappas, I.O., Giannakos, M.N., Krogstie, J., Lekakos, G.: Big data and strategy: a research framework. In: *MCIS*, p. 50 (2016)
8. Brown, B., Chui, M., Manyika, J.: Are you ready for the era of 'big data'. *McKinsey Q.* **4**, 24–35 (2011)
9. Mikalef, P., Pappas, I.O., Krogstie, J., Giannakos, M.: Big data analytics capabilities: a systematic literature review and research agenda. *Inf. Syst. e-Bus. Manage.*, 1–32 (2017)
10. Chen, H., Chiang, R.H., Storey, V.C.: Business intelligence and analytics: from big data to big impact. *MIS Q.* **36**, 1165–1188 (2012)
11. Wang, G., Gunasekaran, A., Ngai, E.W., Papadopoulos, T.: Big data analytics in logistics and supply chain management: certain investigations for research and applications. *Int. J. Prod. Econ.* **176**, 98–110 (2016)
12. Constantiou, I.D., Kallinikos, J.: New games, new rules: big data and the changing context of strategy. *J. Inf. Technol.* **30**, 44–57 (2015)
13. Liu, Y.: Big data and predictive business analytics. *J. Bus. Forecasting* **33**, 40 (2014)
14. Ransbotham, S., Kiron, D.: Analytics as a source of business innovation. *MIT Sloan Manag. Rev.* (2017)
15. Janssen, M., van der Voort, H., Wahyudi, A.: Factors influencing big data decision-making quality. *J. Bus. Res.* **70**, 338–345 (2017)
16. Mikalef, P., Pateli, A.: Information technology-enabled dynamic capabilities and their indirect effect on competitive performance: findings from PLS-SEM and fsQCA. *J. Bus. Res.* **70**, 1–16 (2017)
17. Karahanna, E., Straub, D.W., Chervany, N.L.: Information technology adoption across time: a cross-sectional comparison of pre-adoption and post-adoption beliefs. *MIS Q.* **23**(2), 183–213 (1999)
18. Besson, P., Rowe, F.: Strategizing information systems-enabled organizational transformation: a transdisciplinary review and new directions. *J. Strateg. Inf. Syst.* **21**, 103–124 (2012)
19. Haag, S.: Organizational inertia as barrier to firms' IT adoption—multidimensional scale development and validation (2014)

20. Polites, G.L., Karahanna, E.: Shackled to the status quo: the inhibiting effects of incumbent system habit, switching costs, and inertia on new system acceptance. *MIS Q.* **36**(1), 21–42 (2012)
21. Kelly, D., Amburgey, T.L.: Organizational inertia and momentum: a dynamic model of strategic change. *Acad. Manag. J.* **34**, 591–612 (1991)
22. Hannan, M.T., Freeman, J.: Structural inertia and organizational change. *Am. Sociol. Rev.* **49**, 149–164 (1984)
23. Stieglitz, N., Knudsen, T., Becker, M.C.: Adaptation and inertia in dynamic environments. *Strateg. Manag. J.* **37**, 1854–1864 (2016)
24. Barnett, W.P., Pontikes, E.G.: The Red Queen, success bias, and organizational inertia. *Manag. Sci.* **54**, 1237–1251 (2008)
25. Kim, H.-W., Kankanhalli, A.: Investigating user resistance to information systems implementation: a status quo bias perspective. *MIS Q.* **33**(3), 567–582 (2009)
26. Lyytinen, K., Newman, M.: Explaining information systems change: a punctuated socio-technical change model. *Eur. J. Inf. Syst.* **17**, 589–613 (2008)
27. Rowe, F., Besson, P., Hemon, A.: Socio-technical inertia, dynamic capabilities and environmental uncertainty: senior management view and implications for organizational transformation (2017)
28. Mikalef, P., Framnes, V.A., Danielsen, F., Krogstie, J., Olsen, D.H.: Big data analytics capability: antecedents and business value. In: Pacific Asia Conference on Information Systems (2017)
29. Mikalef, P., Krogstie, J., van de Wetering, R., Pappas, I., Giannakos, M.: Information Governance in the big data era: aligning organizational capabilities. In: Proceedings of the 51st Hawaii International Conference on System Sciences (2018)
30. Benbasat, I., Goldstein, D.K., Mead, M.: The case research strategy in studies of information systems. *MIS Q.* **11**, 369–386 (1987)
31. Battistella, C., De Toni, A.F., De Zan, G., Pessot, E.: Cultivating business model agility through focused capabilities: a multiple case study. *J. Bus. Res.* **73**, 65–82 (2017)
32. Gregor, S.: The nature of theory in information systems. *MIS Q.* **30**, 611–642 (2006)
33. Yin, R.K.: *Case Study Research and Applications: Design and Methods*. Sage Publications (2017)
34. Boudreau, M.-C., Gefen, D., Straub, D.W.: Validation in information systems research: a state-of-the-art assessment. *MIS Q.* **25**(1), 1–16 (2001)
35. Myers, M.D., Newman, M.: The qualitative interview in IS research: examining the craft. *Inf. Organ.* **17**, 2–26 (2007)



Amalgamation of 3D Printing Technology and the Digitalized Industry – Development and Evaluation of an Open Innovation Business Process Model

Danielle Warnecke^(✉), Gor Davidovic Gevorkjan, and Frank Teuteberg

School of Business Administration and Economics, Accounting and Information Systems,
University of Osnabrück, Katharinenstr. 1, 49074 Osnabrück, Germany
{danielle.warnecke, ggevorkjan, frank.teuteberg}@uni-osnabrueck.de

Abstract. Innovations and trends have a significant impact on industries and the way that businesses are conducted. Therefore successful business models need to be increasingly adaptable. We develop a business process model (BPM) focusing primarily on manufacturing companies. This model not only illustrates how the 3D printing (3DP) technology could be implemented in production processes, but also how the approach of open innovation (OI) could be adopted with the purpose of increasing innovative capabilities. We first perform a literature review of the relevant fields; in the second step we develop a BPM (using BPMN 2.0), which is evaluated through expert interviews in a third step. It is demonstrated that the convergence towards highly flexible and customer-oriented production processes and the adoption of OI enhance the amalgamation of 3DP with the digitalized industry.

Keywords: 3D printing technology · Open innovation · Business process model
Digitalization · BPMN 2.0

1 Introduction

According to [1], in the future the manufacturing industry will be reshaped in large parts. This means that advantages of volume production will gradually become insignificant, making traditional production methods obsolete in the long run. 3D printing is mentioned as one of the key trends significantly affecting the manufacturing industry. The range of printable materials continues to expand, while prices continue to fall, making production at low costs economically viable even for standardized parts in the near future [2]. These developments increasingly require manufacturers to consider how advanced information and communication technologies (ICT) might facilitate both higher efficiency and adaptability of their operational processes to ultimately contribute to their value generation [3].

This paper contributes to the ongoing restructuring of the manufacturing industry by developing specific guidance for manufacturers with the purpose of enabling, through

the concept of idea-to-manufacture, a slow transition to an innovative manufacturing organization [4].

[5] also substantiate the importance of OI platforms and 3DP, additionally mentioning the evolvement of a new type of customer who prefers an open, transparent and collaborative environment. A crucial part in developing such a dynamic model is the implementation of an OI approach by illustrating how future manufacturing companies might achieve a smooth transformation into a platform company which leads to the underlying research question: How can open innovation platforms enhance the amalgamation of 3DP and the digitalized industry?

2 Research Methodology

2.1 Research Design

The underlying research process can be summarized in five steps as shown in Fig. 1:




Fig. 1. Research process adapted from [6].

These steps were complemented by the literature review process following [7]. After the broad topic was selected an initial literature review confirmed the existence of a research gap, which specified our research focus. The following search string was developed and applied to the databases AISEL, EBSCO Host, science direct, Springer and web of science:

“open innovation” OR “open business model”) AND (“3D printing” OR “additive manufacturing” OR “rapid prototyping”) AND (“digitaliz*” OR “industry”) AND (“process” OR “crowdsourcing” OR “co-creation”).

In a second step the identified literature was analyzed and filtered for keywords relevant to our research objectives, including business process model (BPM), crucial success factor (CSF) or relevant actors and departments in keywords, titles and abstracts of the literature sample. The third step was to assess the full texts of the extracted literature. The data was considered adequate if it bridges the gap at least in parts and answers important questions within this research. Further it helps to build the fundament in order to establish the pursued business process model. Development of the BPM forms the core of this contribution (Sect. 3). Visualization of the BPM was executed using BPMN 2.0. Additionally, hypotheses were developed in order to evaluate the developed BPM (Sect. 4). For this purpose, seven industry experts were interviewed using a questionnaire that was developed following [8]. The fourth step includes the discussion and interpretation of the findings, which is given in Sect. 5. Step five finally aims to summarize the results, to enable an outlook for the future manufacturing industry and to give directions for further research (Sect. 6).

2.2 Related Work

Industrial digitalization is described as the digital transformation of a company's complete value chain leading to higher automation, sophisticated production and at the same time reaching higher capabilities to control procedures within a corporation and they increase the productivity, efficiency and flexibility [9]. In this context, the development of embedded systems is dominating technology because of its impact on industrial automation and smart production, which in turn, on the bases of new computing infrastructure gradually increases the production performance [10]. The identified publications address the need for transformable business models and their requirements on a theoretical level, but no concrete examples are given or modelled.

However, there is a consensus that highly individualized products of smart manufacturing solutions with real-time connectivity, the integration of business partners and customers into increasingly flexible and interconnected business processes form essential elements of this approach [3, 9, 11].

Thus, a greater product customization seems mandatory, since customers can directly influence the outcome to meet their needs [5], while custom-tailored products lead to a higher degree of differentiation from competition [1].

To meet these requirements 3D printing technology, also known as additive manufacturing or layer-by-layer production, is regarded as one of the disruptive and promising technologies in manufacturing [12], especially concerning faster designs, developing processes and rapid prototyping for components with a complex geometry [13, 14]. 3DP technology moreover allows a faster adaptation to changes in manufacturer's individual needs and market requirements with a more favorable cost structure which matches the asked criteria of smart manufacturing [2, 15]. Hence, our contribution first combines the flexibility of 3DP, a transformable business model and the integration of external knowledge from partners or customers in the innovation process.

In order to enhance the company's capabilities and consequently create value, internal and external know-how can be combined via open innovation (OI). Two different types of OI mechanisms are established, which are described as outside-in and inside-out. "Outside-in" refers to allowing external know-how to access internal barriers and influence internal processes, as in the following BPM [16]. [17] substantiates this idea and characterizes a flexible business model as a crucial part of the innovation process while the highest degree of a flexible business model is a platform that encourages third parties pursuing the same interests to build on the same ideas.

The CSF and corresponding hypotheses for an innovative BPM that meets the requirements of smart manufacturing through the amalgamation of 3DP and the digitalized industry can therefore be derived from the related work, as shown in the following Table 1:

Table 1. Overview of the derived CSF and the resulting hypotheses.

Sources	Derived crucial success factors	Resulting hypotheses
[1, 2, 5, 9, 13, 15, 17]	Customization, agility, mobility of workforce and the adaption of 3DP as a highly flexible technology	H1: The higher the degree of adaptation of 3DP in an organization, the more agile and adjustable are production processes
[3, 5, 18–21]	Integrating tacit knowledge, including individual needs, Co-Creation	H2: The better the integration of consumers in the product development process, the more accurate is the satisfaction of needs
[5, 10, 17, 19, 22]	Quality standards, embedded systems, reliable applications and transparency (to increase productivity)	H3: The fewer barriers are established in an organization, the more likely the OI approach is to succeed
[17, 23, 24]	BPM provides value, success, is adaptable and focuses on customers	H4: The more central the role of the BPM in an organization, the higher the business success
[4, 25, 26]	Platform strategy that exploits open innovation by advanced ICT (online platform, 3DP)	H5: The more an organization transforms into a platform organization, the more likely is the quality of cooperation to decrease

3 Developing a BPM for Future Manufacturers

Our aim is to develop an innovative BPM focusing on the manufacturing industry and to demonstrate how a future manufacturer could implement 3DP as well as the OI approach in their business. To recognize and absorb the dynamics of the environment through an innovative BPM it needs to be well defined [5]. As the use of information systems is highly conceptual and design-oriented, it helps to identify, formalize and visualize the key elements including their relations and operational sequences [27]. To highlight the agile and automatable processes that constitute our BPM we are using BPMN 2.0. The CSF obtained from the literature review function as a starting point for development (see Table 1).

In the following example, three players are included, i.e. the company, a platform and to begin with only one customer. The communicational path between this customer and the platform exists during the entire process and is illustrated as information flows between the pools. In a next step it is aspired to extend the BPM, e.g. to implement multiple customers or partners in the innovation process more deeply.

In this initial version, there are several tasks in the BPM marked as subtasks, which are actually more complex than displayed and would have required an extension of the process. Company A is a manufacturing company and divided into five departments [28]: Management, Technical Department (TD), Marketing Department (MD), Investment Department (ID) and Production Department (ProdD), whereas TD is further divided into Technical Service (TS) and Product Design (PD) and ProdD is further divided into Production Planning (PP) and Production/3DP. The purpose of this classification is to increase the transparency and structure of the example. The following

example displays a complete process, initialized by a customer’s request on the OI platform and finalized with the production process (see Appendix A), described successively.

Step 1 (Fig. 2)

OI-Platform: In order to collaborate with manufacturing company A, a customer submits a detailed description of their product idea after joining the OI platform. On the OI platform all submitted ideas are first collected and allocated before being passed on to either the upper or lower process chain.

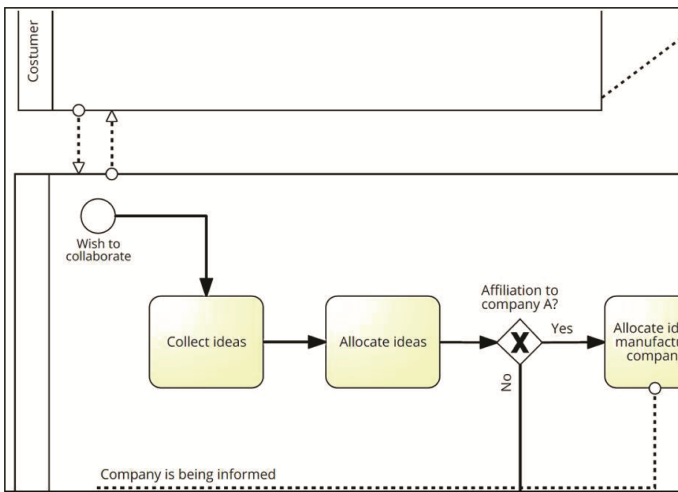


Fig. 2. BPMN Detail 1 - Initializing of cooperation.

Step 2 (Fig. 3)

OI-Platform: In this case, the idea is allocated to manufacturing company A and the manufacturer is informed, so that collaboration with all participants can begin. Additionally, the idea is categorized and passed on to the assessment.

Management: After the company is informed, all information is gathered and the coordination of the involved departments begins. Then the idea is internally classified. After a clear coordination with the Technical Department (TD), the Management awaits the blueprint.

Technical Service (TS): Now the processes in this subdivision begin. The TS checks general technical requirements and provides necessary information for assessment of the idea. Next, the TS coordinates the personnel including company experts.

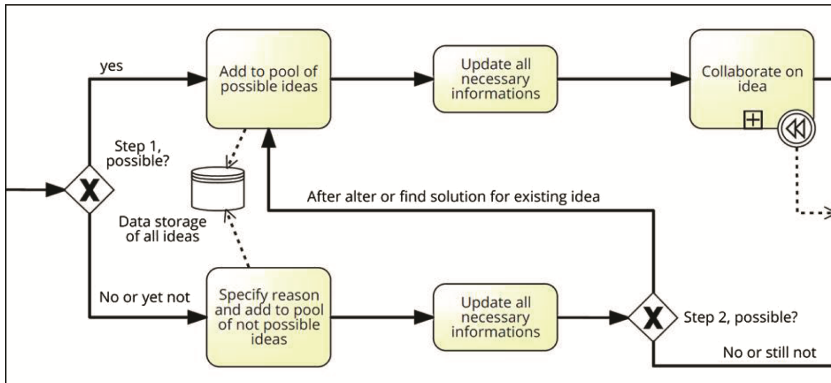


Fig. 3. BPMN Detail 2 – Assessment of ideas.

Step 3

OI-Platform: If the assessment concludes the idea is not feasible, reasons for this will be specified and the idea will be added to the pool of “not possible ideas”. Initially rejected ideas might still be considered at a later stage. If this happened in the BPM, the token would move directly from the gateway “step 1, possible?” via the sequence flow “No or not yet” to the second gateway “step 2, possible?”. This means, that if the idea is not or not yet possible, all necessary information will be updated and coordinated with all parties involved before the project ends. In turn, if the idea is deemed feasible, it will be added to the pool of possible ideas and all necessary information will be updated before the collaboration continues. To handle the possibility that insolvable problems may also occur at a later stage, such as during the collaboration, a compensatory intermediate event is implemented in the BPM, which leads the idea to “not possible ideas”. Ultimately, in a positive case a detailed concept is created (intermediate event) and the OI platform/customer awaits the final product.

Step 4 (Fig. 4)

Product Design (PD): The processes start after the detailed concept is received. In order to develop a design and a blueprint, design developers are provided. The process ends with finalization of the blueprint.

Production Planning (PP): The realization processes start after the blueprint in the PD is available and the coordination has taken place. Next, the request for the PP is checked and aligned with own systems. Schedule time, capacity and other production requirements are checked, so that a needs-based planning can continue. As can be seen in the BPM, another compensatory intermediate event is implemented at this stage in case a needs-based planning is not possible, for instance due to time restrictions or to avoid obstruction of the company’s objectives. Such a compensatory event would lead to an update of all information related to the order, inform the Management, and finally halt the processes. Assuming that all requirements are fulfilled and Management gives their

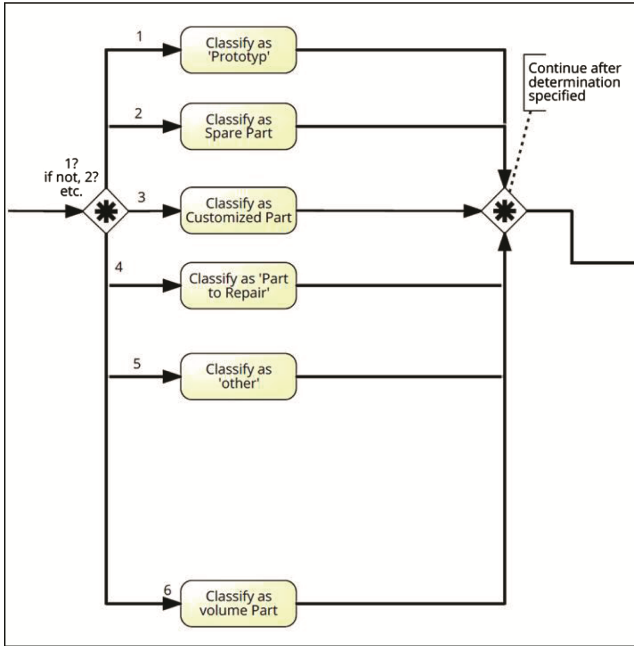


Fig. 4. BPMN Detail 3 – Classification of ideas.

authorization, the PP continues by specifying the order, so that the production process can be determined and the order transmitted appropriately to Production/3DP. For this purpose, the next process step involves a complex gateway allowing an iterative identification of the appropriate production type. The following step is an exclusive gateway. It illustrates that the first five production types are to be produced by 3DP, whereas the fifth type is a dummy position for another classification leading to 3DP and the sixth type is classified as volume production. In this example, the product to be produced is a highly customized product and passed on to the 3DP subdivision. Now, the processes in the PP end.

Management: After authorizing the PP to continue with determining the production type, the Management considers if the product has the potential to contribute to own organizational objectives. Since, in this example, the product idea is too specific and highly customized, the product does not fulfill the requirements for commercialization and so is not included in their own product portfolio.

Marketing Department (MD): Hence, this example does not include criteria for a marketing plan.

Investment Department (ID): The processes begin upon coordination with the Management. As soon as all information required is collected, the evaluation, including a profitability, required budget and prioritization check, begins.

Management: The Management now possesses all information required for a final approval of the product.

Production/3DP: The processes for the final steps are initiated. As already determined, 3DP is required in this example; hence the process chain above the inclusive gateway is activated. The order is now printed and then Management is informed that the product has been completed.

Management: Management finalizes the project.

Step 5 (Fig. 5)

OI-Platform: Once the product has been finalized, the step “Coordination with all parties involved” ensures that all participants receive a progress update before the project is declared successful.

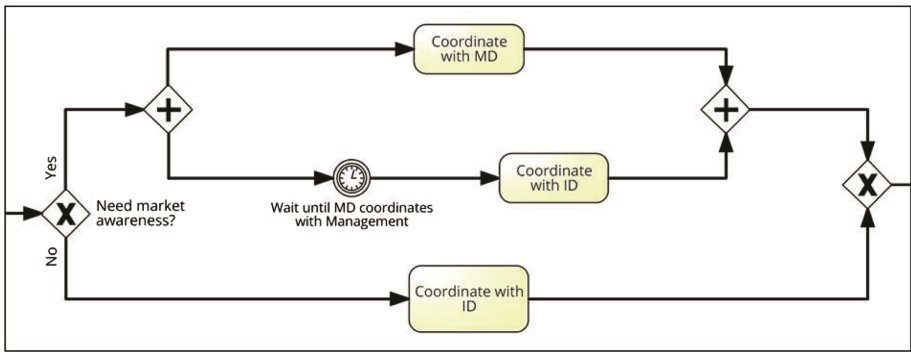


Fig. 5. BPMN Detail 4 – Completion of the process.

4 Evaluation of the BPM

The success of the developed concept depends on the CSF resulting from the literature review (see Table 1). In order to test the five developed hypotheses, a survey was conducted using a questionnaire following [8] (see Appendix B). The interviews were carried out with seven experts working for different internationally active consulting companies in the field of ICT. Divided into sections that aim the five hypotheses the questionnaire consists of 16 questions using a seven step likert scale, which include subquestions at certain points and an optional part for additional information and clarifications. The business process model was also provided as a DIN A2 sized poster and explained briefly. The potential feasibility and practical relevance of the business process model was addressed as final part of the interviews (see Sect. 5).

5 Discussion of Findings

All participants substantiate the indispensability of 3DP for future production processes. They state that spare parts, prototyping, complex parts and final products might be the most attractive applications of 3DP. I4 states: “Trends such as 3DP offer great opportunities for investing into companies’ future growth”, thus underpinning the significance of such an implementation. On the other hand, I2 and I3 state that the significance of 3DP is essential mainly for larger organizations (average value is “slightly disagree”). Nonetheless, all participants agree (average value is 2.25) that the overall benefits outweigh the disadvantages of 3DP (Question 7). Accordingly, H1 can be confirmed.

Regarding H2, I1 indicates: “OI platforms generate more knowledge, they enable more market participants to be included as co-creators and allow the faster detection of changes in customer demands.” The average answer level (2,6) underpins this statement. Additionally, H3 can be confirmed with an answer average score of 2,5. I3 specifies: “System stability plays a crucial role when it comes to the compatibility between new and existing systems [...] Although it is important to consider security aspects, this might prevent the open innovation process from fully unfolding.” Scores for the answers concerning H4 bring up similar conditions (2,6). Given the fact that processes in organizations are becoming more and more digitalized and automated, future BPM need to be increasingly coordinated with the digital world, state I1 and I5.

As an exception, H5 cannot be confirmed. I1 evaluated this aspect with “slightly agree (1)” and added that the technical possibilities to collaborate digitally improve continuously, making physical proximity and its effect on working results lose in significance. However, I3, I6 and I7 evaluated this aspect with “strongly agree (3)”. I3 stated that there are other qualities bound to physical proximity, such as sympathy and direct eye or physical contact, which might influence the collaboration and consequently the working results. Conclusively, a balance between automatized platform organizations and personal interaction must be maintained.

I2 states that the idea of the BPM is plausible in theory but the culture of the company has to be compatible in order to allow the concept to fully unfold its effect. I5 also substantiated that the amalgamation of all mentioned aspects should be feasible, but adds that the willingness of companies to adapt to such a concept largely depends on the perceived additional value. Not only is the relevance of the generated knowledge essential, but also whether this knowledge contributes to the company’s objectives. Overall, the findings substantiate that the implementation of an OI platform and the integration of 3DP will enable more flexible and customer-oriented production processes with the ability to satisfy customer needs to a higher degree and enhance the aforementioned amalgamation (see Table 2).

Table 2. Evaluation results represented as implications of key findings.

Theoretical basis	Implications of key findings
H1: confirmed	<ul style="list-style-type: none"> • 3DP implementation is indispensable for future production of spare parts, prototyping and highly customized products • It might be more relevant for larger companies because of the financial risks
H2: confirmed	<ul style="list-style-type: none"> • OI platforms generate important knowledge from customers and key partners to meet the specific market requirements more effectively
H3: confirmed	<ul style="list-style-type: none"> • System stability and compatibility play a crucial role in transforming and should be regarded at an early stage of the change process • Rigid security policies might decrease the innovation potential
H4: confirmed	<ul style="list-style-type: none"> • Highly digitalized and automated business processes meet the specific needs of markets with highly customized products, especially if these processes are real-time coordinated
H5: not confirmed	<ul style="list-style-type: none"> • Physical proximity, such as sympathy and direct eye contact are still considered very important • Culture and strategy of the transforming company have to be compatible to the platform concept to generate benefits • Willingness to adapt to the concept still depends largely on the expected additional value
Overall	<ul style="list-style-type: none"> • The BPM is plausible and meets the indicated requirements • To represent individual requirements of each transforming company precisely the developed BPM has to be modified

6 Concluding Remark

This contribution considered how the open innovation approach might enhance the amalgamation of the 3D printing technology with the digitalized industry. A BPM was developed and modeled in BPMN 2.0. The result serves as an initial solution for producers to cope with the increasing digitalization of the industry and changing market conditions. A main finding is that future production processes as well as end-user requirements are presumably becoming more complex, raising demand for more adjustable production processes.

However, in order to allow the OI approach to fully uncover its potential, barriers such as certain security aspects must be minimized. Further, it is important to have adjustable business models, enabling the detection and incorporation of essential environmental dynamics. It was indicated that 3DP manufacturing companies may increase their flexibility and enable a stronger customer-orientation by implementing an OI platform. It would be useful to execute a case study with a transforming company to iteratively improve the BPM, including its business processes and practical relevance, and allow insights regarding the quality of the generated knowledge. Moreover, the BPM could be extended to a more interactive version, where multiple customers and partners are enabled to cooperate within the innovation process, including several iteration steps.

Further research could also reveal more information regarding the negative effects of collaborating through platforms rather than in physical proximity, specifically

assessing the paradox mentioned by [29], i.e. exploring whether such strategies contribute to a homogenization of companies, leading them to gradually lose unique features and making it more difficult to achieve a competitive advantage.

Acknowledgments. This work is part of the project “Sustainable Consumption of Information and Communication Technologies in the Digital Society – Dialogue and Transformation through open innovation”. The project is funded by the Ministry for Science and Culture of Lower Saxony and the Volkswagen Foundation (VolkswagenStiftung) through the “Niedersächsisches Vorab” grant programme (grant number VWZN3037).

Appendix: BPM and Questionnaire

Please retrieve Appendix online:

<http://bit.ly/amalgamation-bpmn>

<http://bit.ly/amalgamation-questionnaire>

References

1. D’Aveni, R.: The 3D printing revolution. *Harvard Bus. Rev.* **93**(5), 40–48 (2015)
2. Khan, O., Mohr, S.: 3D printing and its disruptive impacts on supply chains of the future. *Logistics Transport Focus* **18**, 24–28 (2016)
3. Guzman, J.: Digitalization on the horizon. *SMT Mag.*, 62–65 (2015)
4. Johnsson, H.: Production strategies for pre-engineering in house - building: exploring product development platforms. *Constr. Manag. Econ.* **31**(9), 941–958 (2013)
5. Pisano, P., Pironti, M., Rieple, A.: Identify innovative business models: can innovative business models enable players to react to ongoing or unpredictable trends? *Entrepreneurship Res. J.* **5**(3), 181–199 (2015)
6. Baumeister, R., Leary, M.: Writing narrative literature reviews. *Rev. General Psychol.* **1**(3), 311–320 (1997)
7. vom Brocke, J., Simons, A., Riemer, K., Niehaves, B., Plattfaut, R.: Standing on the shoulders of giants: challenges and recommendations of literature search in information systems research. *Commun. Assoc. Inf. Syst.* **37**(1), 205–224 (2015)
8. Gläser, J., Laudel, G.: *Experteninterviews und qualitative Inhaltsanalyse*, 4th edn. VS Verlag, Wiesbaden (2010)
9. Bertschek, I., Clement, R., Buhr, D., Hirsch-Kreinsen, H., Falck, O., Heimisch, A., et al.: *Industrie 4.0: Digitale Wirtschaft - Herausforderung und Chance für Unternehmen und Arbeitswelt*. ifo Schnelldienst, pp. 3–18 (2015)
10. Loupis, M.: Embedded systems development tools: a modus-oriented market overview. *Bus. Syst. Res.* **5**(1), 6–20 (2013)
11. Reischauer, G.: *Industrie 4.0 organisieren*. *zfo Zeitschrift Führung und Organisation*, pp. 271–277 (2015)
12. Palmquist, D.: 5 trends driving the movement toward smart manufacturing. *Supply Demand Chain Executive* **15**, 20–21 (2014)
13. Berman, B.: 3-D printing: The new industrial revolution. *Bus. Horiz.* **55**, 155–162 (2012)
14. Petrick, I., Simpson, T.: *3D Printing Disrupts Manufacturing*. *Research-Technology Management* (2013)

15. Losaw, J.: 3D Printers. *Inventors' Digest*, pp. 30–32 (2016)
16. Chesbrough, H., Bogers, M.: Explicating open innovation: clarifying an emerging paradigm for understanding innovation. In: Chesbrough, H., Vanhaverbeke, W., West, J.: *New Frontiers in Open Innovation*, pp. 3–29. Oxford University Press, Oxford (2014)
17. Chesbrough, H.: Open innovation. *Res. Technol. Manage.* **55**(4), 20–27 (2012)
18. Hansen, M., von Oetinger, B.: Introducing T-shaped managers. *Harvard Bus. Rev.* **79**(3), 106–116 (2001)
19. Christensen, K.: Q&A: the man who wrote the book on open innovation explains how not to get stuck in the 'commodity trap'. *Rotman Mag.*, 84–86 (2012)
20. Grassmann, O.: *R&D Management*, pp. 223–228 (2006)
21. Kleine, D.: The value of social theories for global computing. *Commun. ACM* **58**(9), 31–33 (2015)
22. West, J., Salter, A., Vanhaverbeke, W., Chesbrough, H.: Open innovation: the next decade. *Res. Policy* **43**, 805–811 (2014)
23. Ovans, A.: *Harvard Business Review* (2015). www.hbr.org, <https://hbr.org/2015/01/what-is-a-business-model>. Accessed 9 Nov 2016
24. Chesbrough, H., Rosenbloom, R.: The role of the business model in capturing value from innovation: evidence from Xerox Corporation's technology spin-off companies. *Ind. Corporate Change* **11**(3), 529–555 (2002)
25. Huber, S.: Methodology and tool support for adaptive idea exploitation in open innovation. *Bus. Inf. Syst. Eng.* **59**(5), 1–15 (2016)
26. Edelman, B.: How to launch your digital platform. *Harvard Bus. Rev.* **93**(4), 90–97 (2015)
27. Osterwalder, A., Pigneur, Y.: Designing business models and similar strategic objects: the contribution of IS. *J. Assoc. Inf. Syst.* **14**, 237–244 (2013)
28. Thommen, J.-P., Achleitner, A.-K.: *Allgemeine Betriebswirtschaftslehre. Umfassende Einführung aus managementorientierter Sicht*. Springer Gabler, Wiesbaden (2003)
29. Wada, T., Ichikohji, T., Ikuine, F.: Platform paradox. *Ann. Bus. Admin. Sci.* **13**, 91–103 (2014)

Process Management



Fast Incremental Conformance Analysis for Interactive Process Discovery

P. M. Dixit^{1,2(✉)}, J. C. A. M. Buijs¹, H. M. W. Verbeek¹,
and W. M. P. van der Aalst³

¹ Eindhoven University of Technology, Eindhoven, The Netherlands
{p.m.dixit,j.c.a.m.buijs,h.m.w.verbeek}@tue.nl

² Philips Research, Eindhoven, The Netherlands

³ RWTH, Aachen, Germany
wvdaalst@pads.rwth-aachen.de

Abstract. Interactive process discovery allows users to specify domain knowledge while discovering process models with the help of event logs. Typically the coherence of an event log and a process model is calculated using conformance analysis. Many state-of-the-art conformance techniques emphasize on the correctness of the results, and hence can be slow, impractical and undesirable in interactive process discovery setting, especially when the process models are complex. In this paper, we present a framework (and its application) to calculate conformance fast enough to guide the user in interactive process discovery. The proposed framework exploits the underlying techniques used for interactive process discovery in order to incrementally update the conformance results. We trade the accuracy of conformance for performance. However, the user is also provided with some diagnostic information, which can be useful for decision making in an interactive process discovery setting. The results show that our approach can be considerably faster than the traditional approaches and hence better suited in an interactive setting.

Keywords: Incremental conformance · Interactive process discovery
Domain knowledge · Process mining

1 Introduction

Process mining is a technique that can be used to analyze process-oriented event data from information systems in order to perform business process intelligence tasks. Process mining includes two important tasks: process discovery and conformance checking. Process discovery techniques aim to discover a process model from an underlying event log. The primary aim of a process discovery technique is to come to a visual representation of a process model using the information from the event log. Even though most of the process discovery techniques are automated, there is also a possibility to involve human-in-the-loop in order to interactively construct a process model. Interactive process discovery techniques

combine the traditional worlds of manual process modeling with data support, thereby allowing a user to add domain knowledge during process discovery.

Conformance checking techniques use a process model and an event log to calculate how well an event log fits the process model and vice versa. Traditionally, the conformance checking techniques are used to perform *post-mortem* analysis of the execution of a process. That is, once a process model is made available, conformance checking techniques determine the goodness-of-fit of the event data and the process model. However, in a user-guided process discovery setting, there is a need for active feedback regarding the goodness of a process model after each user interaction (see Fig. 1). That is, after the user interactively changes a process model by taking into account the domain knowledge, there is a need to quantify how good or bad the change was according to the event data. This task can indeed be performed by conformance analysis, and thereby conformance analysis forms a natural fit with user-guided process discovery, and can be used to provide feedback after each step in an interactive process discovery setting.

In an interactive user-guided discovery setting, it is required to have *fast* feedback depending on the change made to the process model. In order to address this problem, we present a framework of calculating incremental conformance depending on the change made in the process model. Unlike other conformance analysis techniques, the framework discussed in this paper exploits the underlying structure used for interactive process discovery in order to perform fast and approximate conformance analysis. Moreover, we present an application of the framework that shows that even though the conformance results are approximated, they still contain diagnostic information which could provide useful feedback for the user for decision making in interactive process discovery.

The rest of the paper is structured as follows. In Sects. 2 and 3, we discuss the related work from literature and the preliminaries resp. In Sect. 4 we discuss our approach and in Sect. 5 we discuss the implementation and visualization details. In Sect. 6 we evaluate our approach and in Sect. 7 we conclude and provide indicators of future work.

2 Related Work

Conformance techniques relate the behavior of processes in real life as depicted by event logs, with the expected behavior as depicted by procedural or declarative process models. A large number of conformance analysis techniques have been introduced within the field of process mining. The authors of [15] were




Fig. 1. Interactive process discovery setting. During discovery, conformance analysis is required to be computed fast enough for an uninterrupted experience for the user.

among the first one's to devise conformance checking techniques in process mining using the token based replay in Petri nets. In [7], the authors discuss conformance strategies for declarative models. The authors of [3,4] use alignment based strategy for performing conformance checking. Some approaches look at the conformance problem from various other dimensions such as natural language processing, formal methods, real time setting etc. [5,6,17]. Most of these approaches focus on the accuracy of the results, and hence do not emphasize on the performance dimension and thereby are not advisable in an interactive process discovery setting. In [18], the authors discuss strategies for incrementally repairing a prior alignment based on a change in a process tree for improving the performance of the ETM algorithm. However, as the authors noted, after a certain number of repairs, the repaired alignment might be drastically different from the so-called optimal alignment. Thereby, this could lead to highly inaccurate results. Also, the class of models supported by this approach is limited to block structured process models (process trees).

In order to improve the performance of conformance analysis, various divide and conquer techniques were discussed in [1,13,16,19,20]. Typically, the central idea behind these techniques is to decompose a process model into various sub-models based on a specific decomposition strategy, and then to compute alignments on the smaller model (and a corresponding smaller log) and aggregate the information across all models. However, in certain scenarios, the decomposed sub-models may still have large state spaces. Therefore, the conformance calculation time in such scenarios may be similar, or even worse owing to decomposition and aggregation time, compared to the complete model. In all these divide-and-conquer techniques, there is more emphasis on the accuracy of the results. However, in our case, we relax the accuracy of the results to ensure short response times, similar to the technique discussed in [11]. In both [11] and our approach, conformance is calculated and aggregated over all combinations of sets of activities for a given cardinality. The main difference between our approach and [11] is that we inherently exploit the incremental nature of process modeling during interactive discovery. That is, in the case of [11], the conformance for all the combinations of activities after each change in a process model is recalculated. However, in our technique, we *remember* the prior conformance results, and typically recalculate only those conformance results which may have changed depending on the change in the process model. Hence, the proposed approach is much faster, robust and provides diagnostic information useful during interactive process discovery.

3 Preliminaries

In this section, we discuss the relevant preliminaries. *Events* form the basic building blocks of an event log. An event represents an occurrence of an activity. Every event has a timestamp associated with it which indicates the time of occurrence of that activity [2]. A trace is a sequence of events, and an event log is a bag (or multiset) of traces.

Having discussed event logs, we now discuss process models. Process models are graphical constructs used to represent the flow of activities. Typically, a process model has a fixed start and a fixed end point, and a navigation from start to end point of a process model results in a trace. Hence a process model corresponds to a set of traces. Conformance analysis in process mining aims at finding a fit between the traces from an event log and a process model.

We now define two concepts of projecting a log and a process model on a set of activities. An event log can be projected on a set of activities by removing all the events from the event log which are not a part of the projection set. Consider an event log $L = [\langle a, b, c, d, a, e \rangle^{10}, \langle a, b, c, d \rangle^5, \langle b, d, e \rangle^5]$. Then the event log projected on activities $\{a, c\}$ is $L \downarrow^{\{a, c\}} = [\langle a, c, a \rangle^{10}, \langle a, c \rangle^5, \langle \rangle^5]$.

A process model M can be projected on a set of activities A by making all the activities from the process model which are not a part of projection set as *invisible* (depicting no behavior) activities. These *invisible* activities can be further removed (if possible) using language preserving reduction rules. The reduction removes the unnecessary *invisible* activities from the net which do not have any behavioral impact on the net. The projected and reduced model is denoted by $M \downarrow^A$. Due to space limitations, we refer the user to [8, 14] for more details about these reduction rules.

4 Incremental Conformance Framework

In this paper we introduce a framework for enabling fast conformance analysis. The high level overview of our framework is presented in Fig. 2. In order to enable fast conformance analysis, we split the problem into two parts, (i) approximating the conformance value by calculating and aggregating the conformance of projected sets of activities and, (ii) incrementally calculating the conformance.

4.1 Projected Conformance

The first part of our approach is similar to the projected conformance checking [11] approach. Here we exploit the fact that calculating conformance for a smaller process model having a smaller state space and a smaller event log is often faster than calculating conformance for a large complex process model. Let A be all the activities present in a model M and let L be the corresponding event log. For a user defined cardinality k , using A we calculate the set of all possible activity combinations C , where $\forall_{c \in C} |c| = k$. Let $M \downarrow^c$ and $L \downarrow^c$ denote the process model M and event log L projected with the activities from activity combination c ($c \in C$). Let $\mathbf{Q}(M)$ define some function used to quantify the quality of the model M w.r.t. the event log. Then the quality of the overall model $\mathbf{Q}(M)$ is the aggregation of the quality values of all the projected models:
$$\mathbf{Q}(M) = \frac{\sum_{c \in C} \mathbf{Q}(M \downarrow^c)}{|C|}.$$

Therefore, instead of calculating the conformance of one (complex) process model with an event log, conformance of several projected (smaller) process




Fig. 2. Overview of conformance calculations when model M_2 is derived from a model M_1 . Projections of models based on activity combinations within a given cardinality are calculated e.g. P1 and P2. For the activity combinations whose projected behavior does not change in M_2 (compared to M_1), e.g. P1 and P2, the projected conformance need not be recalculated. For all the other activity subsets, e.g. P3, the projected conformance needs to be recalculated or newly calculated.

models with projected event logs are calculated and then aggregated. Here we exploit the fact that calculating conformance for one large process model is often time consuming, and thereby distributing the conformance calculation over several smaller projected process models improves efficiency. Similarly, based on user's preference, minimum (maximum) quality value can also be computed of the overall model, using the minimum (maximum) value of each combination.

4.2 Incremental Conformance

Subsection 4.1 provides a way to calculate fast approximate conformance analysis, given a process model and an event log. However, in an interactive process discovery setting, a process model is expanded incrementally. In the second part of our framework, we make use of this principle in order to incrementally calculate the conformance using the projected conformance of the prior model. Instead of recalculating the projected conformance of all the activity combinations, we calculate the projected conformance of only those activity combinations which are necessary, and re-use the previously computed projected conformance for all the other activity combinations. Before introducing incremental conformance, we introduce the concept of behavioral equivalence in two models. Two models M_1 and M_2 are said to be behaviorally equivalent, represented as $M_1 \approx M_2$ iff all the behavior of M_1 is exhibited by M_2 and vice versa. Similarly, behavioral in-equivalence of two models is denoted by $M_1 \not\approx M_2$. Suppose a model M_{i+1} is interactively derived from a model M_i . Let C_{i+1} and C_i correspond to all the activity combinations of M_{i+1} and M_i , with a chosen cardinality k . Then, we can distinguish two cases:

1. Set of *same* activity combinations $C_S \subseteq C_{i+1}$ whose projected behavior is the same in models M_i and M_{i+1} , that is $\forall_{c \in C_S} M_i \downarrow^c \approx M_{i+1} \downarrow^c$.
2. Set of *different* activity combinations $C_D \subseteq C_{i+1}$ whose projected behavior is different in models M_i and M_{i+1} , that is $\forall_{c \in C_D} M_i \downarrow^c \not\approx M_{i+1} \downarrow^c$.

There is no need to calculate the conformance values for those activity combinations which exhibit the same projected behavior in M_i and M_{i+1} (C_S). However, the activity combinations C_D whose projected behavior is *not* the same, there is a need to calculate the conformance values. Therefore, in an interactive setting, we improve the projected conformance calculation times, by calculating conformance for only the activity combinations from the set C_D . It should be noted that, if a new activity is added to the model interactively, then there will be new activity combinations possible, which would all be a part of C_D . The amount of time needed for calculation of conformance for activity combinations C_S is saved by using an incremental way of calculating the conformance values.

4.3 Application of the Framework

We show the instantiation of the framework using a synthesis rules based interactive process modeling approach. In order to do so, we first introduce free choice workflow (FC WF) nets which are a well-known class of Petri nets to represent processes. A workflow (WF) net is a bi partite graph containing places (circles) and transitions (rectangles), such that a transition can only be connected to a place and vice versa. A WF net is navigated from left-to-right, such that the first transition is always a ‘dummy’ (silent) transition called \top , and the last transition is always a ‘dummy’ (silent) transition called \perp . Transitions in a WF net can be labeled with an activity, or they can be *silent*, i.e. not representing any behavior from the event log. Moreover, a WF net is a FC WF net if for any two places, the set of output transitions of both the places are exactly the same, or completely disjoint (that is, there is no overlap at all). Figure 3a shows an example of a FC WF net. The only possible complete trace of this FC WF net is $\langle a, b, c \rangle$. Note that the transitions \top and \perp are silent, and hence do not form a part of this trace sequence. A short circuited version of a FC WF net can be obtained by merging the places i and o , and renaming it as io . It should be noted that we can easily obtain a FC WF net from a short circuited FC WF net by splitting the place io into places i and o .

Three *rules* are used to interact and synthesize a short circuited FC WF net, such that each rule leads to addition of a new place and/or a new transition in the model. Using the synthesis rule based interactive modeling approach, at each iteration we identify the set of combinations of activities (C_S) whose projected behavior does not change. For all the other activity combinations (C_D), the conformance is recalculated. We describe the incremental way of calculating the change in the models depending on each type of rule. It should be noted that, with the usage of synthesis rules, a process model can only grow.

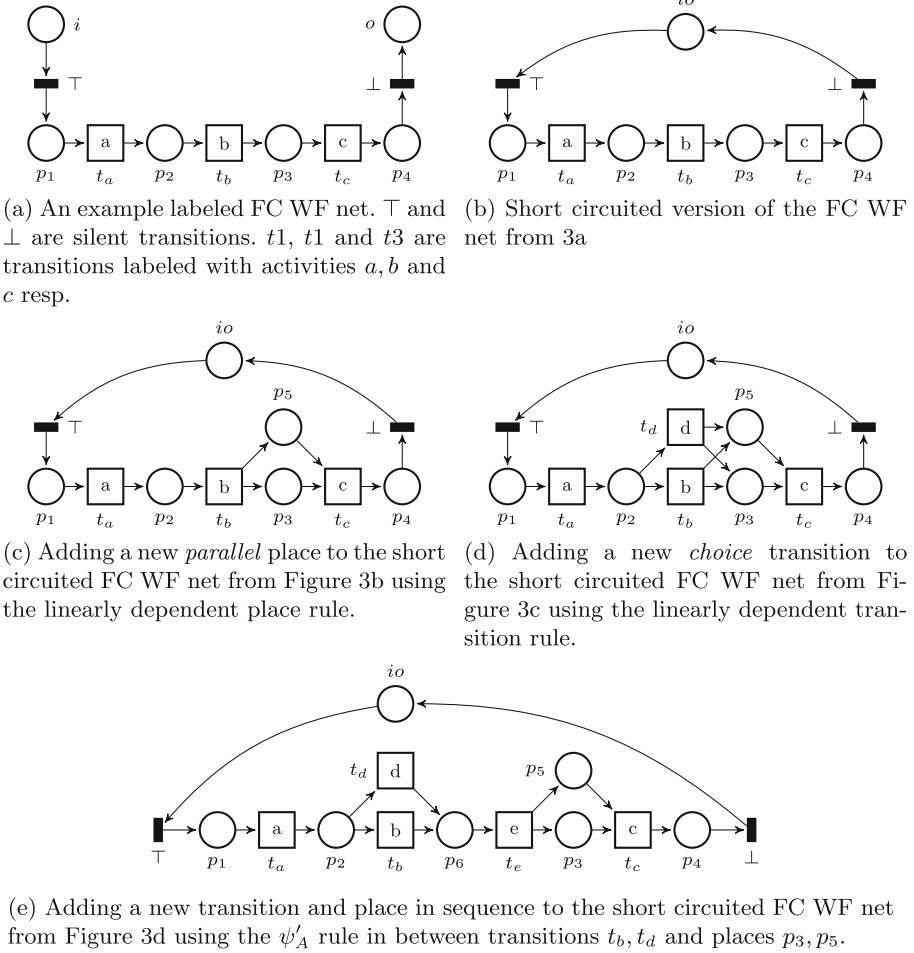


Fig. 3. Synthesis rules [9] applied to short circuited FC WF nets.

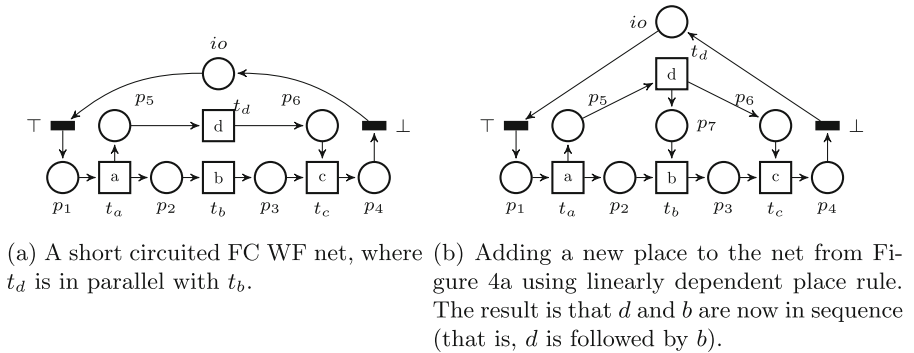



Fig. 4. Addition of a new place resulting in introduction of sequential construct.

4.4 Addition of a New Place

Adding a new place to a net allows introduction of concurrency in the net. An introduction of a place does not result in any new activity in the model, and hence no new activity combinations are possible. There exists a set of bags of places P_{set} in the short circuited FC WF net, which has the same effect as the newly added place. Loosely speaking, this means that every bag of places from P_{set} collectively has the same input and output as the newly added place. In Fig. 3c, this set corresponding to the new place p_5 is $\{[p_3]\}$. Typically, all the activity combinations are added to C_S , as the projected behavior of activity combinations remains unchanged. For example, in Fig. 3c, the projected behavior between activities of the net does not change at all after the addition of the new place. However, in very few cases the projected behavior of activity combinations might have changed, if at least one of the bags in P_{set} contains the place io of the short circuited FC WF net. For example, in Fig. 4b, P_{set} corresponding to the newly added place p_7 is $\{[p_6, p_4, io, p_1, p_2]\}$. Since P_{set} contains io , no activity combinations are added to C_S . This is also because the projected behavior between some activities has indeed changed as shown in Fig. 4 (e.g. t_d and t_b changed from parallel to a sequential construct).

4.5 Addition of a New Transition

Figure 3d shows the addition of a new transition using the so-called linearly dependent transition rule [8]. Addition of a new transition usually results in the introduction of a choice or loop in the model. There exists a set of bags of transitions T_{set} , which have the same effect on the short circuited FC WF net, as the newly added transition. In Fig. 3d, this set corresponding to the newly added transition t_d is $\{[t_b]\}$. We use this information to calculate the set of activity combinations C_S whose projected behavior does not change. No elements are added to C_S , if any bag from T_{set} contains \top or \perp . This is for reasons similar the one described in Subsect. 4.4. The second scenario is when none of the bags from T_{set} contain \top or \perp , i.e. $\forall E \in T_{set} \top, \perp \notin E$. Let T_L be the set of all the labels



(a) Projecting (and reducing) the net from Figure 3d onto $\{b, c\}$. (b) Projecting the net from Figure 3d onto $\{a, c\}$.

Fig. 5. Projected behavior for $k = 2$ corresponding to the new net after adding t_d (Fig. 3d).

represented by the transitions in T_{set} . An activity combination for a subset of activities A_s (s.t. the label of newly added transition is not in A_s) is added to C_S if $A_s \cap T_L = \emptyset$. Consider the model from Fig. 3d derived from Fig. 3c by adding a new transition labeled d ($T_L = \{b\}$). If the cardinality is chosen to be 2, then combination of activities $\{a, b\}$ and $\{b, c\}$ are not added to C_S . For example, consider the projection of activities $\{b, c\}$ as shown in Fig. 5a. In the new net projected (and reduced) on activities $\{b, c\}$, there is a possibility to *skip* the activity b , via τ_d . Hence there is additional behavior introduced corresponding to activity b which was not present in the prior projected net. Hence such activity combinations are not added to C_S , and are candidates for recalculation. As a counter example, it is easy to see that the previous activity pair of $\{a, c\}$ has the same projected behavior, as shown in Fig. 5b, after the introduction of t_d , and hence this activity combination is added to C_S .

4.6 Addition of a New Transition and a New Place

Adding a new transition and a new place using the so-called abstraction rule [8] results in a new sequence in the model. For example, Fig. 3e is derived by adding a new transition (labeled e) and a new place (p_6) to the model from Fig. 3d. If the newly added transition is labeled with an activity which is not already present in the model, then for any chosen cardinality, all the activity combinations from Fig. 3d are behaviorally equivalent in Fig. 3e. That is, if the transition for newly added activity e is made silent, then the net would be behaviorally equivalent to the previous net. Hence all the activity combinations from Fig. 3d are added to C_S .

5 Implementation and Visualization

The technique has been implemented in the Interactive Process Mining package of the nightly build version of ProM tool¹. Figure 6 shows and discusses the visualization of our technique. There are two views, one showing the aggregated visualization of conformance information across all the activity combinations directly on the process model. The other view shows a tabular view of all the activity combinations, along with the corresponding metrics. The user can interact with the activity combinations from the tabular view, and visualize the reduced models containing only those activities present in the selected activity combination. This allows the user to dig deeper to analyze the intricacies of fragmented process models with a certain cardinality, as shown in Fig. 8.

¹ <http://www.promtools.org/doku.php?id=nightly>.

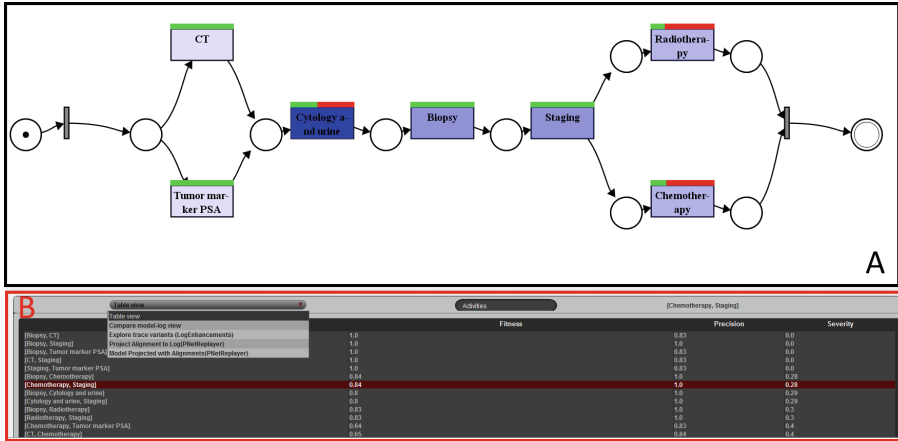


Fig. 6. (A) shows the conformance of the modeled activities according to the event log: the density of blue color indicates the frequency, the red (darker) and green (lighter) bars above the transitions indicate the distribution of fitting and non-fitting events. The tabular view (B) shows the individual scores of the activity sets ($k = 2$ in this case). The user can dig deeper to explore the relationship between any activity set by choosing the desired visualization. (Color figure online)

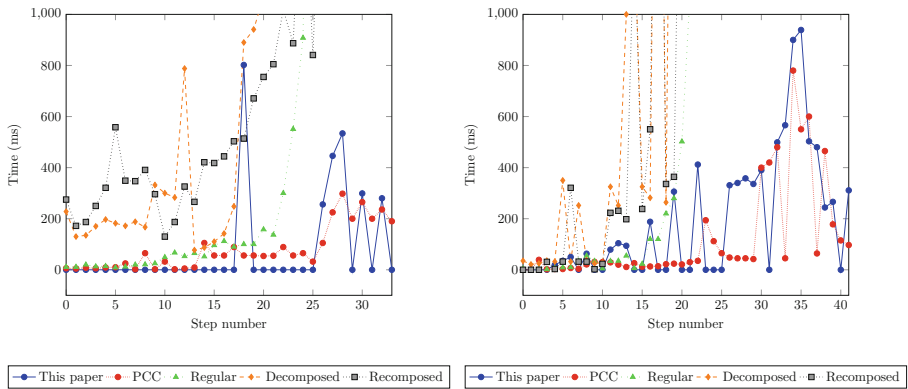
6 Evaluation

We evaluate the approach presented in this paper by comparing it with state-of-the-art conformance techniques. The goal of this evaluation is to show the effectiveness of our approach in an interactive setting, measured in terms of performance times and correctness of the result. We use two real-life event logs: (i) the Sepsis event log² containing the workflow data for roughly 1000 patients suffering from Sepsis in a hospital. A normative process model is already available for this event log at [12]. (ii) the BPIC 2011 event log³ containing hospital data for cancer patients from a Dutch hospital - filtered to retain only top 40% of the most frequent activities, resulting in a total of around 10000 events. We use the inductive miner [10] to discover a process model from the filtered event log. In order to replicate an interactive scenario, starting with a process model without any activities, we interactively re-construct each process model using the synthesis rules based approach. After each “interaction” (step), conformance is recalculated using following techniques: the decomposed replay technique [20], the recomposed replay technique [19], projected conformance checking [11] (with $k = 2$), the technique presented in this paper (with $k = 2$) and the regular alignments technique [4]. The fitness (i.e., the faithfulness of the model to the log), precision (i.e., the extent to which behavior not seen in the log is allowed by the model) and time taken for recalculation w.r.t. each technique are recorded.

² <https://data.4tu.nl/repository/uuid:915d2bfb-7e84-49ad-a286-dc35f063a460>.

³ <http://www.win.tue.nl/bpi/2011/challenge>.

The fitness and precision values are scored from 0 (worst) to 1 (best) in all the techniques.



(a) Performance times for the filtered BPIC 2011 event log.

(b) Performance times for the Sepsis event log.

Fig. 7. The time(s) taken after each interaction (step) for conformance analysis.

Figure 7 compares the performance of each approach in terms of time taken after each step. It is quite evident that the traditional approaches, along with decomposed and recomposed approaches can be extremely slow, especially as the size of the process model increases. It was observed that the traditional approach for calculating fitness/precision could take more than 30 min s for both the event logs for the final step. It should however be observed that even though the decomposed approaches are slower, the quality values computed using these techniques were identical to the alignment based conformance technique [4] (which can be considered as a baseline for fitness value). It should be noted that increasing the value of k can potentially improve the accuracy, however even with a value of $k = 2$ our approach is within 5% of the baseline and is much faster than the traditional approaches as we exploit the inherent rules used during process composition. [11] is typically faster than most of the approaches. However, [11] provides very limited diagnostic information about the transitions in the FC WF net, especially when there are duplicate activities present in the model. That is, if there are duplicate occurrences of an activity in the FC WF net, then [11] would assign them the same fitness and precision value. Our approach is more robust and can distinguish the difference in behavior of duplicate occurrences of activities in the FC WF net, for e.g. Fig. 8. Moreover, by projecting the model on a subset of activities, as shown in Fig. 8, we can find the relationship between activities which may not be directly connected in the overall model. Also, in contrast to all the other techniques, during many steps in the process construction, the response time of our approach is almost zero, or close to zero. These are the

changes in the process model wherein the projected behavior of all the activity combinations was the same in the iterated model.




Fig. 8. The view of projected alignment on activities ‘CRP’ and ‘ER Registration’ from the Sepsis event log.

7 Conclusion and Future Work

In this paper, we presented a framework to enable fast conformance checking in an interactive process discovery setting. We instantiated this framework to exploit the underlying principles used in interactive process discovery to calculate fast conformance by incrementally studying the change in the structure of a model. By using two real-life event logs we were able to show that the approach suggested in this paper is faster compared to many state-of-the-art conformance checking techniques. Furthermore, even though the actual fitness (and precision) scores are only approximated, they are still very close to the original values as computed using the traditional alignment-based conformance checking approach. In the future, we would like to extend the technique presented in this paper to combine it with other conformance checking techniques such as the one in [11]. Furthermore, we would also like to explore the effect of different values of k in terms of performance time and accuracy.

References

1. van der Aalst, W.M.P.: Decomposing petri nets for process mining: a generic approach. *Distrib. Parallel Databases* **31**(4), 471–507 (2013)
2. van der Aalst, W.M.P.: *Process Mining - Data Science in Action*, 2nd edn. Springer, Heidelberg (2016). <https://doi.org/10.1007/978-3-662-49851-4>
3. Adriansyah, A., van Dongen, B.F., van der Aalst, W.M.P.: Conformance checking using cost-based fitness analysis. In: 2011 15th IEEE International Enterprise Distributed Object Computing Conference (EDOC), pp. 55–64. IEEE (2011)
4. Adriansyah, A., van Dongen, B.F., van der Aalst, W.M.P.: Towards robust conformance checking. In: zur Muehlen, M., Su, J. (eds.) *BPM 2010. LNBIP*, vol. 66, pp. 122–133. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-20511-8_11
5. vanden Broucke, S.K.L.M., Munoz-Gama, J., Carmona, J., Baesens, B., Vanthienen, J.: Event-based real-time decomposed conformance analysis. In: Meersman, R., Panetto, H., Dillon, T., Missikoff, M., Liu, L., Pastor, O., Cuzzocrea, A., Sellis, T. (eds.) *OTM 2014. LNCS*, vol. 8841, pp. 345–363. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-662-45563-0_20

6. Carmona, J.: The alignment of formal, structured and unstructured process descriptions. In: van der Aalst, W., Best, E. (eds.) *PETRI NETS 2017*. LNCS, vol. 10258, pp. 3–11. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-57861-3_1
7. de Leoni, M., Maggi, F.M., van der Aalst, W.M.P.: Aligning event logs and declarative process models for conformance checking. In: Barros, A., Gal, A., Kindler, E. (eds.) *BPM 2012*. LNCS, vol. 7481, pp. 82–97. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-32885-5_6
8. Desel, J., Esparza, J.: *Free Choice Petri Nets*, vol. 40. Cambridge University Press, Cambridge (2005)
9. Esparza, J.: Synthesis rules for Petri nets, and how they lead to new results. In: Baeten, J.C.M., Klop, J.W. (eds.) *CONCUR 1990*. LNCS, vol. 458, pp. 182–198. Springer, Heidelberg (1990). <https://doi.org/10.1007/BFb0039060>
10. Leemans, S.J.J., Fahland, D., van der Aalst, W.M.P.: Discovering block-structured process models from event logs containing infrequent behaviour. In: Lohmann, N., Song, M., Wohed, P. (eds.) *BPM 2013*. LNBIP, vol. 171, pp. 66–78. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-06257-0_6
11. Leemans, S.J.J., Fahland, D., van der Aalst, W.M.P.: Scalable process discovery and conformance checking. *Software & Systems Modeling*, July 2016
12. Mannhardt, F., Blinde, D.: Analyzing the trajectories of patients with sepsis using process mining. In: *RADAR+EMISA*, vol. 1859, pp. 72–80 (2017)
13. Munoz-Gama, J., Carmona, J., van der Aalst, W.M.P.: Single-Entry Single-Exit decomposed conformance checking. *Inf. Syst.* **46**, 102–122 (2014)
14. Murata, T.: Petri nets: properties, analysis and applications. *Proc. IEEE* **77**(4), 541–580 (1989)
15. Rozinat, A., van der Aalst, W.M.P.: Conformance checking of processes based on monitoring real behavior. *Inf. Syst.* **33**(1), 64–95 (2008)
16. Taymouri, F., Carmona, J.: Model and event log reductions to boost the computation of alignments (2016)
17. Taymouri, F., Carmona, J.: A recursive paradigm for aligning observed behavior of large structured process models. In: La Rosa, M., Loos, P., Pastor, O. (eds.) *BPM 2016*. LNCS, vol. 9850, pp. 197–214. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-45348-4_12
18. Vázquez-Barreiros, B., van Zelst, S.J., Buijs, J.C.A.M., Lama, M., Mucientes, M.: Repairing alignments: striking the right nerve. In: Schmidt, R., Guédria, W., Bider, I., Guerreiro, S. (eds.) *BPMS/EMMSAD-2016*. LNBIP, vol. 248, pp. 266–281. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-39429-9_17
19. Verbeek, H.M.W.: Decomposed replay using hiding and reduction as abstraction. In: Koutny, M., Kleijn, J., Penczek, W. (eds.) *Transactions on Petri Nets and Other Models of Concurrency XII*. LNCS, vol. 10470, pp. 166–186. Springer, Heidelberg (2017). https://doi.org/10.1007/978-3-662-55862-1_8
20. Verbeek, H.M.W., van der Aalst, W.M.P., Munoz-Gama, J.: Divide and conquer: a tool framework for supporting decomposed discovery in process mining. *Comput. J.* **60**(11), 1649–1674 (2017)



Business Process Compliance and Business Process Change: An Approach to Analyze the Interactions

Tobias Seyffarth^(✉), Stephan Kuehnel, and Stefan Sackmann

Martin Luther University Halle-Wittenberg, 06108 Halle (Saale), Germany
{tobias.seyffarth, stephan.kuehnel,
stefan.sackmann}@wiwi.uni-halle.de

Abstract. The adherence of business process compliance (BPC) is crucial for many companies. In addition, business processes may be supported by IT components, which can also be affected by compliance requirements. Due to business process change and the avoidance of compliance violations, companies must analyze, among other things, the interactions between business process change and BPC. Following the design science research paradigm, we developed and prototypically implemented a method that is able to analyze interactions between BPC and business process change considering supporting IT components and compliance processes. The method takes the business process change patterns “replace” and “delete” into account.

Keywords: Business process compliance · Business process change
Compliance process · Information technology

1 Introduction

The adherence of compliance requirements resulting from the interpretation of laws, standards, or business contracts is crucial for many companies [1, 2]. The execution of business processes in accordance with their applicable compliance requirements is called business process compliance (BPC) [3]. Nevertheless, compliance requirements can affect both business processes and their supporting information technology (IT) [4–7]. However, many factors, such as new technologies, improvement of business processes, and outsourcing decisions can lead to business process changes, including changed IT components or compliance requirements [8, 9]. In dynamic markets, the fast detection of compliance violations and the adherence to the demands of compliance requirements to changed business processes and supporting IT components are necessary [10].

So far, an approach to determine the interactions between BPC and business process change considering supporting IT components automatically is lacking. Existing approaches only solve a small part of the problem. For example, in [11], a language to query business process models is presented. The presented software artifact in [8] automatically detects effects on business activities due to changed compliance requirements. In [12] an approach to model compliance requirements, business processes, and their

supporting IT components together is proposed. Consequently, we address the following research questions (RQ):

- RQ 1: *What are the interactions between BPC and business process change considering supporting IT components?*
- RQ 2: *How can these interactions be determined automatically?*

We address this research questions by developing two different artifacts. In this paper, we focus on their design and development according to the design science research paradigm by Hevner et al. [13]. The artifacts are each an exaptation that extends known solutions to new problems. Following the problem statement, the developed artifacts address a relevant problem and the required scientific rigor is fulfilled by the usage of existing methodologies. With the utilization of methodologies in the field of graph modeling and graph search (e.g. [14]) in addition to already existing software libraries for the development of our software prototype [15, 16], we present the following contributions:

- We propose a method that can analyze the interactions between BPC and business process change considering supporting IT components and compliance processes.
- We propose a prototypical implementation of our method. It uses models based on common standards such as the Business Process Model and Notation (process models), ArchiMate (IT architecture models), and laws that are provided as XML files by the German Federal Ministry of Justice (compliance requirements) [17–19].

This contribution is structured as follows: In Sect. 2, the backgrounds of business process compliance in conjunction with enterprise architecture and business process change patterns is briefly introduced. In Sect. 3, a running scenario is presented to illustrate the problem statement. In Sect. 4, a method that analyzes the interactions between BPC, supporting IT components and business process change is shown. In addition, the prototypical implementation of the method is presented. In Sect. 5 the related work is discussed. Finally, a brief conclusion and research outlook are provided in Sect. 6.

2 Theoretical Background

2.1 Business Process Compliance

As already mentioned, business process compliance (BPC) denotes the execution of business processes in adherence to applicable compliance requirements [3]. In this case, a compliance requirement is a constraint or assertion that results from the interpretation of compliance sources, such as laws, regulations or standards (e.g., [1, 2, 20]).

Various approaches check or ensure BPC. As an example, BPC can be checked after process execution by analyzing log files (e.g., [21]) or ensured at the design time of business processes. Here, a possible solution is the separate modeling of so-called compliance processes and its integration into business processes. Whenever a business process is affected by a compliance requirement, an appropriate compliance process can be integrated in the business process [1, 22]. In this context, a compliance process is

defined as an independent process (part) consisting of at least one compliance-related activity that ensures BPC [23].

2.2 Enterprise Architecture

An architecture is defined as the fundamental organization of a system embodied in its components, their relationships to each other and to the environment, and the principles guiding its design and evolution [24]. Following this, an enterprise architecture is understood as the fundamental organization of a government agency or a corporation, either as a whole or together with partners, suppliers and/or customers (“extended enterprise”), or in part (e.g., a division, a department, etc.) as well as the principles governing its design and evolution [25].

In the literature, various sub-architectures of enterprise architectures have been discussed. For example, Winter and Fischer [26] divided an enterprise architecture into a business, process, integration, software, and technology architecture. The Open Group split an enterprise architecture into the business, data, application and technology architecture [27]. Within business architecture, both Winter and Fischer [26] as well as the Open Group [27] have situated business processes. Within software/application and infrastructure architecture, the application services, application components or network devices as special hardware types are situated. For reasons of simplicity, we refer to a single element within an IT architecture as an IT component. In this paper, we do not distinguish between different types of IT components. Nevertheless, in the following we will focus on three perspectives of an enterprise architecture: (1) business process (2), compliance, and (3) IT architecture. The interrelations between the single elements of these perspectives are illustrated in Fig. 1.




Fig. 1. Considered elements and their interrelations (based on [20, 23])

We focus on these perspectives for many reasons: The perspectives on business processes and compliance are considered because of the definition of BPC. As mentioned earlier, BPC denotes the execution of business processes in adherence to applicable compliance requirements [3]. Furthermore, we split the compliance perspective into the elements compliance requirement and the compliance process. As already mentioned, a compliance process operationalizes the compliance requirement that affects a single business activity [23]. Thus, the compliance process within a business process helps to satisfy the compliance requirement [22, 23]. We consider IT components as single elements within an IT architecture for three reasons. First, IT components are sometimes necessary to execute a business activity. Second, within an automatic

compliance process, IT components are also necessary for the proper execution of automatic compliance processes [23]. Third, both business processes and also IT components are affected by numerous compliance requirements. For example, COSO [7] and the Institute of Public Auditors in Germany [6] provide a lot of requirements for the proper operation of IT components.

2.3 Business Process Change Patterns

In addition to many sub-architectures of an enterprise architecture, various business process change patterns have been discussed, as well. For example, Weber et al. [28] distinguished 18 change patterns and split them into adaption patterns and patterns of change in predefined regions of business processes. As shown in Fig. 2, further authors have combined the proposed 18 changed patterns to four change patterns: “insert element”, “delete element”, “replace element” and “update element” [29, 30].




Fig. 2. Business process change patterns (based on [29, 30])

The insert change pattern inserts a new element into the business process at a defined place. The delete pattern removes an existing element. The replace pattern replaces an existing element with a new one. The update pattern modifies an attribute of an existing element [30].

These change patterns can also be applied to our perspectives on enterprise architecture. Therefore, we assume that IT components, compliance requirements, and compliance processes are additional views of a business process that are similar to for example, a data view. Due to space limitations, we will only apply the change patterns “delete element” and “replace element”. Next, we discuss a running scenario to motivate our research problem and to demonstrate the solution through our method.

3 Running Scenario

The left side of Fig. 3 shows a simplified purchase to pay process based on [31, 32], including perspectives on compliance requirements, compliance processes, and IT components. Some of the business activities are supported by IT components that are modeled as triangles.




Fig. 3. Running scenario (based on [31, 32])

We assume that the material management module of an enterprise resource planning system (ERP MM) is a prerequisite to creating a purchase request. Furthermore, a financial module of an ERP system (ERP FI) is a prerequisite for both the compliance process “approve invoice payment” and the business activity “pay invoice”.

In addition to that, both business activities and IT components are sometimes affected by compliance requirements. In this example, “§ 238 of the German Commercial Code” demands the accounting obligation for merchants in Germany. It is also related to the proper operation of IT that supports the accounting activities. In this example, the compliance requirements “physical access” and “logical access” are prerequisite by “§ 238 of the German Commercial Code”. The compliance requirement “physical access” requires a regulated access to physical IT components while “logical access” requires the identification and authentication of users of an application [33]. Furthermore, the compliance requirement “internal policy” specifies additional requirements that are necessary for the payment of invoices. The compliance process “approve invoice payment” helps to satisfy the “internal policy”.

The right side of Fig. 3 shows (1) the compliance requirements that shall be taken into consideration when replacing the IT component “ERP FI” and (2) the effects on BPC by deleting “ERP FI”. The edge direction corresponds to that of the model on the left side. Our method and prototypical implementation can automatically calculate these interactions considering the business process change patterns “replace element” and “delete element”.

In the case of replacing “ERP FI”, there are two compliance requirements that place direct demands to them. In addition to the compliance requirement “logical access”, the “physical access” compliance requirement must also be taken into account. Furthermore, there is a transitive relation by “internal policy” and “§ 238 of the German Commercial Code”. The IT component “ERP FI” is a prerequisite for the business activity “pay invoice” which is affected by the “internal policy” which is, in turn, the “§ 238 of the German Commercial Code” is the prerequisite of the “internal policy”.

In the case of deleting “ERP FI” the compliance process “approve invoice payment” cannot be executed. Thus, the compliance requirements “internal policy” and “§ 238 of the German Commercial Code” are violated. Furthermore, the compliance requirement “logical access” is obsolete. In the next section, we propose a method and its prototypical implementation for analyzing the interactions between BPC and business process change considering supporting IT components and compliance processes.

4 Interaction Between BPC and Business Process Change

Figure 4 shows the three steps of our method for analyzing the interactions of BPC and business process change. First, business process, compliance process, IT architecture, and compliance requirements must be imported into the software prototype. Second, the imported models must be transformed and integrated into a single graph. Third, the interactions of BPC and business process change must be analyzed. In the following sections, we explain each step in detail with respect to the running scenario introduced in the previous section.




Fig. 4. Steps to analyze the interactions of BPC and business process change

4.1 Import Models

All above mentioned model types must be available as formal models for an import into our prototype. We used the process modeling language Business Process Model and Notation (BPMN) as the import format for both business and compliance process models, which offers various XML schemas for the process model description [17]. The modeling language ArchiMate [18] was used as the import format for IT architectures for two reasons. First, it provides a formal data exchange format [27]. Second, ArchiMate is a promising approach to model an enterprise architecture on the bottom of the BPMN process models [34]. Finally, we used several German laws, which were provided as XML files by the German Federal Ministry of Justice [19] as an import format for compliance requirements.

Our prototype was developed in Java. For the deserialization of the BPMN process according to the XML schemas proposed by the OMG [17], we used the Java library “Camunda BPMN model API” [15]. It allows the deserialization of each element within the process model into a single Java object. For the deserialization of the IT architecture and compliance requirements we each developed a XML wrapper.

4.2 From Separate Models to an Integrated Graph

We adapted graph search methods (e.g., [14]) to analyze the interaction of BPC and business process change considering supporting IT components and compliance processes. In order to formalize this concept, a number of auxiliary concepts must be introduced. We define a non-empty finite set of directed and labeled multigraphs $G = (V, E, F, H, I)$ with its elements $g_i \in G$. In addition, V is a nonempty finite set of vertices, $e_l \in E$ is a set of directed edges between two vertices (v_i, v_j) and $f_i \in F$ is the unique identification (id) of the vertex v_i . In addition, $h_i \in H$ is the model type of the vertex v_i with $H = \{BP, CP, IT, CR\}$, which corresponds to the imported model types business process (BP), compliance process (CP), IT architecture (IT) or compliance requirement (CR). Finally, $i_l \in I$ is the label of edge e_l .

The transformation of the previously imported models to an integrated graph consists of two major steps. First, each imported model is transformed into a graph g_i and added to the graph set G . The transformation of the imported models works as follows: Within business and compliance processes the “Camunda BPMN model API” [15] allows the extraction of the lists of all flow nodes and sequence flows of a process model. Based on the list of flow nodes it is possible to create a vertex $v_{i,n}$ for each flow node. Therefore, each vertex id $f_{i,n}$ is equal to the id of the flow nodes within the process models. The model type $h_{i,n}$ of each vertex is either a business process or compliance process. The list of sequence flows contains the predecessor and successor relations of all flow nodes. Based on this predecessor and successor relations, the appropriate edges e_l between the vertices can be created. The respective edge labels i_l are equal to the edge labels connecting different model types in the legend of Fig. 3. Similarly to business and compliance process models, the imported IT architecture also contains a list of IT components and a list of relations that describes the IT component predecessor and successor relations. Therefore, the graph of the IT architecture can be generated similarly to the graph of a process model. Finally, each compliance requirement and its corresponding compliance source are represented by two vertices that are connected by an edge.

Second, the generated graphs must be integrated into a single graph. Therefore, an empty graph is generated and all previously generated graphs within the graph collection G are added to it. Then, the required edges between the vertices of the business process model, compliance process model, IT components, and compliance requirements must be created manually. However, the edges that are connected to a process model can also be modeled automatically. Whenever a flow node of a business or compliance process is the target vertex of a compliance requirement or an IT component, additional meta-information is added to the respective flow node. This is done by using the BPMN class extension element [17]. The extension is a list of key-value-pairs. Thereby, the keys represent the referenced model types h_i and the values their respective id f_i . Thus, the edges that touch a business or compliance process can also be modelled automatically. This is always possible when the business or compliance process contains an aforementioned BPMN extension element.

Within our software prototype, we used the Java library “JGraphT” [16] to model and visualize graphs. In the next step, the integrated graph is used as a basis for analyzing the interactions between BPC and business process change.

4.3 Analyzing the Interactions Between BPC and Business Process Change

By taking the business process change patterns “replace element” and “delete element” into consideration, their semantic interpretation to the business process model considering IT components, compliance processes, and compliance requirements must be defined. When replacing an element, the replaced elements are sometimes affected by compliance requirements (see replace “ERP FI” within the running scenario). Thus, the new element must take these respective compliance requirements into consideration to ensure BPC. The removal of an element can also have an impact on BPC (see remove “ERP FI” within the running scenario). First, it can have an impact on compliance requirements and/or compliance processes that satisfy the impacted compliance requirements. Second, due to a removal of an element, the respective compliance requirements can become obsolete. In the following, we describe our algorithm to analyze the interactions between BPC and business process change. Due to space limitations, we will refer to an IT component as the replaced or deleted element.

Replacing an Element

Figure 5 shows the algorithm to analyze the relations of compliance requirements and compliance processes to the IT component that shall be replaced. To ensure BPC, the replaced element must fulfill the related compliance requirements and be able to execute a related compliance process. Thereby, we distinguish between direct and transitive relations.

<pre> Input: Graph <i>g</i>, element that shall be replaced $v \in g$ where $h(v)=it$ architecture 1 // get all direct related compliance requirements and compliance processes to <i>v</i> 2 ForEach <i>i</i> in (get all predecessor of <i>v</i> where $h=$compliance requirement or $h=$compliance process) do 3 mark <i>i</i> as direct AND add <i>i</i> including all vertices between <i>i</i> und <i>v</i> to <i>result</i> 4 // get all transitive related compliance processes and compliance requirements to <i>v</i> 5 ForEach <i>it</i> in (get all leafs of <i>v</i> where $h=it$ architecture) do 6 ForEach <i>activity</i> in (get all direct successor of <i>it</i> where $h=$business process) do 7 ForEach <i>cr</i> in (get all direct predecessor of <i>activity</i> where $h=$compliance requirement) do 8 <i>i</i> = get all predecessor of <i>cr</i> 9 mark <i>it</i>, <i>activity</i>, <i>cr</i> and <i>i</i> as transitive AND add to <i>result</i> 10 ForEach <i>complianceprocess</i> in (get all direct successor of <i>it</i> where $h=$compliance process) do 11 ForEach <i>cr</i> in (get all direct successor of <i>complianceprocess</i> where $h=$compliance requirement) do 12 <i>i</i> = get all predecessor of <i>cr</i> 13 mark <i>it</i>, <i>complianceprocess</i>, <i>cr</i> and <i>i</i> as transitive AND add to <i>result</i> 14 generate <i>result_graph</i> based on <i>g</i> and <i>result</i> Output: Graph <i>result_graph</i> </pre>
--

Fig. 5. Algorithm to analyze interactions by replacing an IT component

A direct relation between compliance requirements and a replaced IT component occurs in two different ways. First, each compliance requirement that is directly related to the replaced IT component must be considered. Within our running scenario the compliance requirement “logical access” is directly related to the replaced IT component “ERP FI”. Consequently, the compliance requirement “§ 238 of the German Commercial Code” is also directly related, because it is a prerequisite for “logical access”.

Second, all compliance requirements of the replaced element prerequisites must be taken into consideration, as well. Within the running scenario, the compliance requirement “Physical access” is direct related to the replaced element “ERP FI” because it is a prerequisite for the IT component “hardware”, which is, in turn a prerequisite for “ERP FI”.

A transitive relation between an updated IT component and a compliance requirement or a compliance process also occurs in two different cases. First, the changed IT component is a prerequisite of a business activity that is affected by at least one compliance requirement. Second, the updated IT component is a prerequisite of a compliance process that helps to satisfy at least one compliance requirement. Within the running scenario, this is the case at the business activity “pay invoice” and its prerequisite compliance requirements “internal policy” and “§ 238 of the German Commercial Code”.

Deleting an Element

Figure 6 shows the algorithm to analyze the impacts on BPC by removing an IT component. Here, we distinguish between two different impacts. First, the removal an IT component can lead to a compliance process not being executable. This concerns the compliance process “approve invoice payment”, since the IT component “ERP FI” is a prerequisite for its execution. Consequently, the compliance requirements “internal policy” and “§ 238 of the German Commercial Code” are violated.

<pre> Input: Graph g, element that shall be removed v ∈ g where h(v)=it architecture 1 // get all violated compliance requirements 2 Foreach it in (get all leaves of v where h=it architecture) do 3 Foreach complianceprocess in (get all direct successor of it where h=compliance process) do 4 Foreach cr in (get all direct successor of complianceprocess where h=compliance requirement) do 5 i = get all predecessor of cr 6 mark cr and i as violated AND add cr and i to result 7 // get all obsolete compliance requirements 8 add v to list_it 9 Foreach it in (get all successor of v where h=it architecture) do 10 If (it only requires v) 11 add it to list_it 12 Foreach it in (list_it) do 13 Foreach cr in (get all direct predecessor of it where h=compliance requirement) do 14 If (get all direct successor of cr=it) 15 mark cr as obsolete AND add cr to result 16 Foreach cr2 in (get all predecessor of cr where h=compliance requirement) do 17 If (cr2 hasn't other direct successor than get all predecessor of cr where h=comp.req.) 18 mark cr2 as obsolete AND add cr2 to result 19 generate result_graph based on g and result Output: Graph result_graph </pre>

Fig. 6. Algorithm to analyze interactions by deleting an IT component

Second, the removal of an IT component can may cause a compliance requirement to become obsolete within the respective model. This concerns the compliance requirement “logical access”. It only affects the removed IT component “ERP FI”. Nevertheless, its prerequisite compliance requirement “§ 238 of the German Commercial Code” does not become obsolete. It is already a prerequisite for the compliance requirements “physical access” and “internal policy”.

5 Related Work

We conducted a structured literature review according to vom Brocke et al. [35] to retrieve the related work. As proper databases we used AISel, IEEE, SpringerLink as well as the EbscoHost libraries “Science & Technology Abstracts”, “Business Source Premier” and “Academic Source Premier”. With respect to our research question, we defined the following abstract query: << (“business process compliance” OR BPC) AND (“information technology” OR flexibility OR query OR change)>>. Furthermore we performed a backward search according to Webster and Watson [36]. These searches identified a total of 552 unique hits. After analyzing the articles abstract, keywords and full text, if necessary, we excluded 519 articles. In total, 33 articles were read in full.

Finally, the articles can be classified into two main topics. There are articles that discuss (1) querying of business processes and (2) the linkage and (partly) the querying of the considered enterprise architecture perspectives on business processes, compliance, and IT architecture. Nevertheless, none of the proposed approaches analyzed the interaction between BPC and business process change, considering IT components or compliance processes. Due to space limitations, we will highlight the main ideas of the selected approaches.

Awad [11] proposed a visual language called BPMN-Q to query BPMN process models. Delfmann et al. [37] developed a generic model query language (GMQL). GMQL is able to express pattern queries for all graph-based modelling languages by treating the modeling language as an attributed graph. In addition to structural aspects of graph matching, Gacitua-Decar and Pahl [38] proposed an approach to identify patterns in business processes based on structural and semantic aspects. Fellmann et al. [39] transformed a business process model into an ontology and enriched the ontology instance with structural and domain representation information; they queried business processes by using SPARQL queries within the ontology.

Several authors have proposed approaches to link compliance requirements with business processes, such as Ghanavati et al. [40]. Rudzajs and Buksa [8] developed an approach to detect changes in compliance requirements by storing them into a version control system. Thus, it is able to detect the business activities affected by the changed compliance requirements. Furthermore, both Fdhila et al. [41] and Knuplesch et al. [42] analyzed the possible effects on both compliance requirements and business processes in the context of collaborative business processes.

Only a few authors have discussed the link between compliance requirements and IT components, such as Knackstedt et al. [5]. Furthermore, Knackstedt et al. [43] extended their previous work by developing a software prototype that models and shows the link between compliance requirements and IT components.

In addition, only a few authors have discussed the linkage between different compliance requirements. Independently of each other, Elgammal et al. [44] and Halle [45] presented approaches to detect root cause violations in rules that are formalized in linear temporal logic.

Furthermore, Koetter et al. [12, 46] presented an approach to link compliance requirements to both business processes and IT components. They used so-called compliance descriptors to formalize compliance requirements for business processes

and IT components. The formalized compliance requirements allow the determination of the interrelation between business processes or IT components and compliance requirements.

6 Conclusion and Outlook

New technologies, improvement of business processes, and outsourcing decisions can lead to changing business processes, IT components, or compliance requirements [8, 9]. Therefore, the fast detection of compliance violations and the adherence to demands of compliance requirements to changed business processes and IT components are necessary [10]. To solve that challenge, we proposed a method to analyze the interactions between BPC and business process change considering IT components and compliance processes. Hence, we focused on the perspectives business processes, compliance processes, compliance requirements, and IT components of an enterprise architecture. Furthermore, we took the business process change patterns “replace element” and “delete element” into consideration. On the one hand, a replaced element can be affected directly or transitive by compliance requirements. On the other hand, the removal of an element can lead to a violation or an irrelevance of compliance requirements. In addition, a prototype implementation of the method was shown. The prototype is based on well-established libraries for handling BPMN and graphs in Java [15, 16].

In the first step, we only considered the aforementioned perspectives of an enterprise architecture. Unquestionably, there are more perspectives, such as data or an organizational perspective (e.g., [26, 27]). Apart from that, there are also more preconditions than the existence of IT components, e.g., data or organizational units, to execute a compliance process [23]. Thus, we plan to extend our graph search algorithm and its demonstration to other perspectives of an enterprise architecture. In addition, the applicability of the proposed method will be evaluated. Another open research question is the adaption of a business process in case a compliance process is not executable. In this case, the integration of an alternative compliance process seems to be a promising approach to ensure BPC despite a business process change. Furthermore, we plan to extend our prototype for an automatic integration of alternative compliance processes.

References

1. Schumm, D., Turetken, O., Kokash, N., Elgammal, A., Leymann, F., van den Heuvel, W.-J.: Business process compliance through reusable units of compliant processes. In: Daniel, F., Facca, F.M. (eds.) ICWE 2010. LNCS, vol. 6385, pp. 325–337. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-16985-4_29
2. Turetken, O., Elgammal, A., van den Heuvel, W.-J., Papazoglou, M.: Enforcing compliance on business processes through the use of patterns. In: 19th ECIS 2011 (2011)
3. Schäfer, T., Fettke, P., Loos, P.: Towards an integration of GRC and BPM – requirements changes for compliance management caused by externally induced complexity drivers. In: Daniel, F., Barkaoui, K., Dustdar, S. (eds.) BPM 2011. LNBIP, vol. 100, pp. 344–355. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-28115-0_33

4. Sadiq, S., Governatori, G., Namiri, K.: Modeling control objectives for business process compliance. In: Alonso, G., Dadam, P., Rosemann, M. (eds.) BPM 2007. LNCS, vol. 4714, pp. 149–164. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-75183-0_12
5. Knackstedt, R., Eggert, M., Heddier, M., Chasin, F., Becker, J.: The relationship of is and law - the perspective of and implications for IS research. In: ECIS 2013 Completed Research (2013)
6. The Audit of Financial Statements in an Information Technology Environment. IDW AuS 330 (2002)
7. Committee of Sponsoring Organizations of the Treadway Commission (COSO): Internal Control - Integrated Framework. Framework and Appendices (2012)
8. Rudzajs, P., Buksa, I.: Business process and regulations: approach to linkage and change management. In: Grabis, J., Kirikova, M. (eds.) BIR 2011. LNBIP, vol. 90, pp. 96–109. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-24511-4_8
9. Fdhila, W., Indiono, C., Rinderle-Ma, S., Reichert, M.: Dealing with change in process choreographies: design and implementation of propagation algorithms. *Inf. Syst.* **49**, 1–24 (2015)
10. Rinderle, S., Reichert, M., Dadam, P.: Correctness criteria for dynamic changes in workflow systems—a survey. *Data Knowl. Eng.* **50**, 9–34 (2004)
11. Awad, A.: BPMN-Q: a language to query business processes. In: Proceedings of EMISA 2007, pp. 115–128 (2007)
12. Koetter, F., Kochanowski, M., Weisbecker, A., Fehling, C., Leymann, F.: Integrating compliance requirements across business and IT. In: 2014 IEEE 18th International Enterprise Distributed Object Computing Conference (2014)
13. Hevner, A.R., Gregor, S.: Positioning and presenting design science research for maximum impact. *MIS Q.* **37** (2013)
14. Yu, P.S., Han, J., Faloutsos, C. (eds.): Link Mining: Models, Algorithms, and Applications. Springer Science + Business Media LLC, New York (2010). <https://doi.org/10.1007/978-1-4419-6515-8>
15. Camunda: Camunda BPMN model API. <https://github.com/camunda/camunda-bpmn-model>
16. Naveh, B.: JGraphT. <http://jgrapht.org/>
17. OMG (Hg): Business Process Model and Notation (BPMN). <http://www.omg.org/spec/BPMN/2.0/PDF/>
18. Jonkers, H., Lankhorst, M., van Buuren, R., Hoppenbrouwers, S., Bonsangue, M., van der Torre, L.: Concepts for modeling enterprise architectures. *Int. J. Coop. Inf. Syst.* **13**, 257–287 (2004)
19. Juris (ed.): Gesetze im Internet. <http://www.gesetze-im-internet.de>
20. Seyffarth, T., Kühnel, S., Sackmann, S.: ConFlex: an ontology-based approach for the flexible integration of controls into business processes. *Multikonferenz Wirtschaftsinformatik (MKWI) 2016*, 1341–1352 (2016)
21. Kharbili, M., Medeiros, A.K.A.d., Stein, S., van der Aalst, W.M.P.: Business process compliance checking: current state and future challenges. *MobIS* **141**, 107–113 (2008)
22. Sackmann, S., Kittel, K.: Flexible workflows and compliance: a solvable contradiction?! In: Vom Brocke, J., Schmiedel, T. (eds.) BPM - Driving Innovation in a Digital World, pp. 247–258. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-14430-6_16
23. Seyffarth, T., Kühnel, S., Sackmann, S.: A taxonomy of compliance processes for business process compliance. In: Carmona, J., Engels, G., Kumar, A. (eds.) BPM 2017. LNBIP, vol. 297, pp. 71–87. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-65015-9_5
24. IEEE: IEEE Recommended Practice for Architectural Description of Software Intensive Systems, (IEEE Std 1 1471–2000). IEEE Computer Society, New York (2000)

25. TOGAF (ed.): Content Metamodel. Content Metamodel Vision and Concepts. <http://pubs.opengroup.org/architecture/togaf9-doc/arch/>
26. Winter, R., Fischer, R.: Essential layers, artifacts, and dependencies of enterprise architecture. In: 2006 10th IEEE International Enterprise Distributed Computing Conference Workshops (EDOCW 2006), p. 30 (2006)
27. The Open Group (ed.): TOGAF 9.1. Content Meta Model. <http://pubs.opengroup.org/architecture/togaf9-doc/arch/chap34.html>
28. Weber, B., Reichert, M., Rinderle-Ma, S.: Change patterns and change support features – enhancing flexibility in process-aware information systems. *Data Knowl. Eng.* **66**, 438–466 (2008)
29. Rinderle-Ma, S., Reichert, M., Weber, B.: On the formal semantics of change patterns in process-aware information systems. In: Li, Q., Spaccapietra, S., Yu, E., Olivé, A. (eds.) ER 2008. LNCS, vol. 5231, pp. 279–293. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-87877-3_21
30. Fdhila, W., Rinderle-Ma, S., Reichert, M.: Change propagation in collaborative processes scenarios. In: 8th International Conference on Collaborative Computing: Networking, Applications and Worksharing (CollaborateCom) (2012)
31. Namiri, K.: Model-Driven Management of Internal Controls for Business Process Compliance. Karlsruhe (2008)
32. Frank, U., Heise, D., Ulrich, Kattenstroth, H., Ferguson, D.F., Hadar, E., Waschke, M.G.: ITML: a domain-specific modeling language for supporting business driven IT management. In: Proceedings of the 9th OOPSLA Workshop on Domain-Specific Modeling (2009)
33. Principles of Proper Accounting When Using Information Technology. IDW AcP FAIT 1 (2002)
34. Kirikova, M., Penicina, L., Gaidukovs, A.: Ontology based linkage between enterprise architecture, processes, and time. *Commun. Comput. Inf. Sci.* **539**, 382–391 (2015)
35. Vom Brocke, J., Simons, A., Niehaves, B., Riemer, K., Plattfaut, R., Cleven, A.: Reconstructing the giant: on the importance of rigour in documenting the literature search process. In: 17th European Conference on Information Systems, pp. 2206–2217 (2009)
36. Webster, J., Watson, R.T.: Analyzing the past to prepare for the future: writing a literature review. *MIS Q.* **26**, xiii–xxiii (2002)
37. Delfmann, P., Steinhorst, M., Dietrich, H.-A., Becker, J.: The generic model query language GMQL – conceptual specification, implementation, and runtime evaluation. *Inf. Syst.* **47**, 129–177 (2015)
38. Gacitua-Decar, V., Pahl, C.: Automatic business process pattern matching for enterprise services design. In: 2009 World Conference on Services - II (2009)
39. Fellmann, M., Thomas, O., Busch, B.: A query-driven approach for checking the semantic correctness of ontology-based process representations. In: Abramowicz, W. (ed.) BIS 2011. LNBIP, vol. 87, pp. 62–73. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-21863-7_6
40. Ghanavati, S., Amyot, D., Peyton, L.: Compliance analysis based on a goal-oriented requirement language evaluation methodology. In: 2009 17th IEEE International Requirements Engineering Conference (2009)
41. Fdhila, W., Rinderle-Ma, S., Knuplesch, D., Reichert, M.: Change and compliance in collaborative processes. In: 2015 IEEE International Conference on Services Computing (2015)

42. Knuplesch, D., Fdhila, W., Reichert, M., Rinderle-Ma, S.: Detecting the effects of changes on the compliance of cross-organizational business processes. In: Johannesson, P., Lee, M.L., Liddle, Stephen W., Opdahl, Andreas L., López, Ó.P. (eds.) ER 2015. LNCS, vol. 9381, pp. 94–107. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-25264-3_7
43. Knackstedt, R., Braeuer, S., Heddier, M., Becker, J.: Integrating regulatory requirements into information systems design and implementation. In: ICIS 2014 Proceedings (2014)
44. Elgammal, A., Turetken, O., van den Heuvel, W.-J., Papazoglou, M.: Root-cause analysis of design-time compliance violations on the basis of property patterns. In: Maglio, Paul P., Weske, M., Yang, J., Fantinato, M. (eds.) ICSSOC 2010. LNCS, vol. 6470, pp. 17–31. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-17358-5_2
45. Halle, S.: Causality in message-based contract violations. a temporal logic “Whodunit”. In: 2011 IEEE 15th International Enterprise Distributed Object Computing Conference (2011)
46. Koetter, F., et al.: An universal approach for compliance management using compliance descriptors. In: Helfert, M., Ferguson, D., Méndez Muñoz, V., Cardoso, J. (eds.) CLOSER 2016. CCIS, vol. 740, pp. 209–231. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-62594-2_11



Mining Hybrid Business Process Models: A Quest for Better Precision

Dennis M. M. Schunselaar^{1(✉)}, Tijs Slaats², Fabrizio M. Maggi³,
Hajo A. Reijers¹, and Wil M. P. van der Aalst⁴

¹ Vrije Universiteit Amsterdam, Amsterdam, The Netherlands
{d.m.m.schunselaar,h.a.reijers}@vu.nl

² University of Copenhagen, Copenhagen, Denmark
slaats@di.ku.dk

³ University of Tartu, Tartu, Estonia
f.m.maggi@ut.ee

⁴ RWTH Aachen University, Aachen, Germany
wvdaalst@pads.rwth-aachen.de

Abstract. In this paper, we present a technique for the discovery of hybrid process models that combine imperative and declarative constructs. In particular, we first employ the popular Inductive Miner to generate a fully imperative model from a log. Like most imperative miners, the Inductive Miner tends to return so-called flower models for the less structured parts of the process. These parts are often imprecise. To counter these imprecise parts, we replace them with declarative models to increase the precision since declarative models are good at specifying which behavior is disallowed. The approach has been implemented in ProM and tested on several synthetic and real-life event logs. Our experiments show that hybrid models can be found to be more precise without overfitting the data.

Keywords: Hybrid process model · Process mining
Process discovery · Process tree · Declare

1 Introduction

In recent years, different comparative investigations have been conducted to better understand the distinctive characteristics of imperative and declarative process modeling languages and to support the choice of the most suitable paradigm to be used in different scenarios [25, 26]. While advantages and limitations of the two paradigms are still a matter of investigation, both in academic research and in industry [12, 21], a trend has emerged to consider hybrid approaches combining a mixture of imperative and declarative specifications. The motivations behind this trend rely on the surmise that many real-life processes are characterized by

This work is supported in part by the Hybrid Business Process Management Technologies project funded by the Danish Council for Independent Research.

a mixture of (i) less structured processes with a high level of variability, which can usually be described in a compact way using declarative languages such as Declare [24], SCIFF [23], or DCR Graphs [10], and (ii) more stable processes with well-structured control flows, which are more appropriate for traditional imperative languages such as Petri nets and BPMN.

As a result, there have been several recent efforts to fully develop and formalize such hybrid modeling notations and methodologies. These hybrid approaches can be categorized based on how tightly the imperative and declarative paradigms are integrated, with each approach falling in one of three major categories: (i) a *complementary* approach, where imperative and declarative notations are used in separate models, without any semantic interaction, but with the aim of representing different aspects of the same process (see, e.g., [11]). (ii) A *hierarchical* approach, where both imperative and declarative notations are used in the same overall model, but the model is separated into several sub-processes and each sub-process uses a single uniform notation. To our knowledge most existing proposals for hybrid approaches fall within this category (see, e.g., [1, 27, 30]). (iii) A fully *mixed* approach, where imperative and declarative constructs fully overlap in the same model. While this approach can be considered the most flexible, it is also the most complex since the modeler needs to consider and understand in minute detail how the elements of the different notations interact. An example of such an approach is the one proposed in [36].

One field where the hybrid approach may lead to significant advances is that of automatic process discovery [5], which aims at generating useful process models from real-life event logs. Many imperative and declarative process discovery algorithms (usually referred to as miners) exist, and it remains an open challenge to determine a suitable discovery approach given the characteristics of the event logs [7]. When faced with very variable logs, imperative miners have a tendency to create “spaghetti models”, or “flower models”. The former are generally unreadable, the latter allow any behavior and therefore contain very little useful information about the process. Declarative miners, on the other hand, are not well-suited to dealing with very structured logs: when faced with such a log they will often generate an abundance of constraints, which makes the model unreadable for human users. Hybrid process discovery approaches aim at providing a middle way between these two extremes, generating an imperative model for the structured parts of the log and declarative constraints for the unstructured ones. Two primary approaches to hybrid discovery have been considered: (i) a *model-driven* approach, where first an imperative or a declarative model is created, from which “bad” parts are identified (e.g., flower sub-models in imperative models, or over-constrained activities in declarative models) and secondly these “bad” parts are re-mined with a more suitable mining algorithm. (ii) A *log-driven* approach where an up-front analysis of the log is used to identify structured and unstructured parts, which are then separated into sub-logs. Each sub-log is mined with the most suitable discovery algorithm.

We can categorize the full spectrum of hybrid miners by orthogonally considering their overall mining technique and the type of models that they output.

The current paper falls within the Hierarchical/Model-driven section of this categorization. In particular, we propose a hybrid miner that builds on the Inductive Miner [16] by Leemans et al., which, for a given input log, generates an imperative process model in the form of a process tree. The Inductive miner guarantees a fitness of 1. This guarantee sometimes comes at the cost of precision, e.g., when no structure can be found in the log, the Inductive miner defaults to a flower loop, which is able to capture all behavior and much more. Being hierarchical in nature, process trees can be easily adopted into a hierarchical hybrid approach. To do so, we take the process tree returned from the Inductive miner and identify nodes from the process tree to be replaced with Declare models. We have used Declare since there are several well accepted mining algorithms for this notation, whereas support to process discovery for other declarative notations is either non-existent or in a very early stage.

In this paper, we primarily focus on the precision of the model. There are two main reasons for this: (1) precision, contrary to other quality dimensions like simplicity, has widely accepted metrics for measuring it, and (2) flower loops tend to be very imprecise since they allow for any possible behavior, thus giving an opportunity to have a more accurate understanding of the process captured in the event log by replacing them with declarative models.

The approach has been implemented as a plug-in of the process mining tool ProM and tested on both synthetic and real-life logs. The evaluation shows positive results for both synthetic and real-life logs, with up to 52.18% improvement in the precision of the mined models for synthetic logs and up to 3471% increase in precision for the real-life logs.

The rest of the paper is structured as follows: in Sect. 2, we discuss the related work. In Sect. 3, we introduce the hybrid models we use. In Sect. 4, we describe the proposed approach. In Sect. 5, we report on our evaluation. Section 6 concludes the paper and spells out directions for future work.

2 Related Work

In the literature, there are different families of approaches that have been proposed for the discovery of imperative process models. The first family of approaches extracts some footprint from the event log and uses this footprint to directly construct a process model (see, e.g., [5, 34]). The second family of process discovery approaches first constructs an intermediate model and then converts it into a more refined model that can express concurrency and other advanced control-flow patterns (e.g., [4]). The third family of approaches tries to break the discovery problem into smaller problems. For example, the Inductive miner [16] aims at splitting the event log recursively into sub-logs, which are eventually mined. Techniques originating from the field of computational intelligence form the basis for the fourth family of process discovery approaches. An example of this type of approaches is the genetic mining approach described in [22].

In the area of declarative process discovery, Maggi et al. [18, 19] first proposed an unsupervised algorithm for mining Declare models. They base the discovery

of constraints on the replay of the log on specific automata, each accepting only those traces that are compliant with one constraint. In [8, 14, 15], the authors use inductive logic programming techniques for the supervised learning of constraints expressed using SCIFF [6]. The approach for the discovery of Declare constraints presented in [13] is divided in two steps. The first step computes statistic data describing the occurrences of activities and their interplay in the log. The second step checks the validity of Declare constraints by querying the statistic data structure built in the previous step. In [28], the authors present a mining approach that works with RelationalXES, a relational database architecture for storing event log data. The relational event data is queried with conventional SQL.

Recent research has put into evidence synergies between imperative and declarative approaches [25, 26]. Accordingly, hybrid process modeling notations have been proposed. In particular, in [9], the authors provide a conservative extension of BPMN for declarative process modeling, namely BPMN-D, and show that Declare models can be transformed into readable BPMN-D models. In [36], the authors propose to extend Colored Petri nets with the possibility of linking transitions to Declare constraints directly. The notion of *transition enablement* is extended to handle declarative links between transitions. A recent implementation of this technique has been made available in CPN Tools 4.0 [35]. In [31], the authors extend the work in [36] by defining a semantics based on mapping Declare constraints to R/I-nets and by proposing modeling guidelines for the mixed paradigm. In [32], a discovery approach based on this paradigm is proposed. Differently from these approaches devoted to obtain a fully mixed language, in [30], a hybrid process model is hierarchical, where each of its sub-processes may be specified in either a procedural or declarative fashion. In [20], an approach for the discovery of hybrid process models based on this semantics is presented. In contrast to the approach of the current paper, which is model-driven, the approach in [20] is log-driven. In addition, our approach guarantees that the resulting hybrid model has 100% fitness.

3 Hybrid Models

As discussed in the introduction, we use the hierarchical approach to hybrid models. In our previous work [30], we have formalized the execution semantics of any combination of different formalisms used in a hierarchical hybrid model, e.g., Petri nets, BPMN, Declare, DCR Graphs, and process trees. In the mining approach presented here, we limit ourselves to hybrid models consisting of process trees and Declare models, and, in particular, to hybrid models where the top-level model is a process tree and some sub-processes can be Declare models.

In Fig. 1, we show an example of a hybrid model. The top-level model is a process tree, and there are two declarative sub-processes. The top-level process tree has the regular activities a , b , c , d , e and o and two abstract activities $X1$ and $X2$, each mapping to a declarative sub-process. The declarative sub-process $X1$ contains activity s and the sub-process $X2$ contains activities p , q , r and t .




Fig. 1. Example of a Hybrid Model

3.1 Process Trees

Process trees are a tree-structured modeling formalism. The branches of the trees represent control-flow constructs and the leaves of the tree represent atomic activities (e.g., a), including silent activities (τ). Various dialects exist for process trees (see, e.g., the ones used in [16, 29]). We adopt the one from [16], which has four types of control-flow nodes: (i) the sequential composition (\rightarrow) denotes that the child nodes should occur in sequence, (ii) the exclusive choice (\times) denotes that (exactly) one of the child nodes should occur, (iii) the parallel composition ($+$) denotes that the child nodes should occur in parallel, and (iv) the redo loop (\odot), which contains a *do* part (the first child) and a *redo* part (all subsequent children), and denotes a kind of looping behavior where the *do* is the main body of the loop and one can repeat the loop after doing one of the *redo* children.

We can define the semantics of process trees in terms of the formal languages (sets of sequences) that they represent. We reuse the definition from [2], but extend it to handle sub-process nodes. In the definition, we use two special operators on languages: (i) the concatenation operator (\cdot) concatenates the languages, e.g., $\{\langle a \rangle\} \cdot \{\langle b \rangle\} \cdot \{\langle c \rangle, \langle d \rangle\} = \{\langle a, b, c \rangle, \langle a, b, d \rangle\}$, and (ii) the shuffle operator (\diamond) returns all possible interleavings of the languages, e.g., $\{\langle a \rangle\} \diamond \{\langle b, c \rangle, \langle d, e \rangle\} = \{\langle a, b, c \rangle, \langle b, a, c \rangle, \langle b, c, a \rangle, \langle a, d, e \rangle, \langle d, a, e \rangle, \langle d, e, a \rangle\}$.

Definition 1 (Process Tree Semantics). For a process tree Q with alphabet A , the language $\mathcal{L}(Q)$ of the process tree is defined recursively over its nodes as:

1. $\mathcal{L}(\tau) = \{\langle \rangle\}$;
2. $\mathcal{L}(a) = \{\langle a \rangle\}$, if $a \in A$;
3. $\mathcal{L}(\overline{X}) = \mathcal{L}(\mathcal{M}_X)$, where \mathcal{M}_X is the model of the sub-process X ;
4. $\mathcal{L}(\rightarrow (Q_1, Q_2, \dots, Q_n)) = \mathcal{L}(Q_1) \cdot \mathcal{L}(Q_2) \cdot \dots \cdot \mathcal{L}(Q_n)$;
5. $\mathcal{L}(\times (Q_1, Q_2, \dots, Q_n)) = \mathcal{L}(Q_1) \cup \mathcal{L}(Q_2) \cup \dots \cup \mathcal{L}(Q_n)$;
6. $\mathcal{L}(+(Q_1, Q_2, \dots, Q_n)) = \mathcal{L}(Q_1) \diamond \mathcal{L}(Q_2) \diamond \dots \diamond \mathcal{L}(Q_n)$;
7. $\mathcal{L}(\odot (Q_1, Q_2, \dots, Q_n)) = \{\sigma_1 \cdot \sigma'_1 \cdot \sigma_2 \cdot \sigma'_2 \cdot \dots \cdot \sigma_m \in A^* \mid m \geq 1 \wedge \forall_{1 \leq j \leq m} \sigma_j \in \mathcal{L}(Q_1) \wedge \forall_{1 \leq j < m} \sigma'_j \in \cup_{2 \leq i \leq n} \mathcal{L}(Q_i)\}$.

Table 1. Semantics of Declare templates

Template	Semantics	Notation
existence(a)	a must occur at least once	
absence2(a)	a can occur at most once	
init(a)	a must occur in the first position	
responded existence(a, b)	If a occurs, b must occur as well	
response(a, b)	If a occurs, b must eventually follow	
alternate response(a, b)	If a occurs, b must eventually follow, without any other a in between	
chain response(a, b)	If a occurs, b must occur next	
precedence(a, b)	b can occur only if a has occurred before	
alternate precedence(a, b)	b can occur only if a has occurred before, without any other b in between	
chain precedence(a, b)	b can occur only immediately after a	
not chain succession(a, b)	If a occurs, b cannot occur next	
not succession(a, b)	If a occurs, b cannot eventually follow	
not co-existence(a, b)	If a occurs, b cannot occur	

Given this semantics, the language of the process tree in Fig. 1 is: $\{\langle a, b, X1, X2, o, c, d, e \rangle, \langle a, b, X2, X1, o, c, d, e \rangle, \langle a, b, X2, o, X1, c, d, e \rangle, \langle a, b, X1, X2, c, d, e \rangle, \langle a, b, X2, X1, c, d, e \rangle\}$.

3.2 Declare

Declare [24] is a declarative process modeling language that encodes a set of declarative templates into a formal logic. Encodings into several different logics have been proposed, such as a linear temporal logic (*LTL*) and regular expressions. When using Declare in the context of process mining, it is most sensible to use one of their finite variants, as event logs contain only finite traces. Therefore, we employ the encoding into the finite variant of *LTL* (*LTL_f*) for the current paper. Table 1 shows an overview of the core relations used in Declare models. For the *existence* and *absence* templates, the language supports all countable variations, e.g., 0..2, 0..3, 4..*, etc.

Definition 2 (Declare Model). A Declare model is a pair $M = (A, C)$, where A is a set of activities and C is a set of constraints over the activities in A .

The language of a Declare model $M = (A, C)$ is the set of all traces that satisfy every constraint in C . Formally:

Definition 3 (Declare Semantics). The language $\mathcal{L}(M)$ of a Declare model $M = (A, C)$ is defined as: $\mathcal{L}(M) = \{\sigma \in A^* \mid \forall_{c \in C} \sigma \models c\}$.




Fig. 2. Overview of the Approach

Sub-processes $X1$ and $X2$ in Fig. 1 are examples of Declare models. $X1$ is very simple and contains three constraints: $absence3(s)$, $existence2(s)$ and $init(s)$, indicated by the two boxes on activity s , one containing the text 2..2, and the other the text $init$. These constraints together indicate that s is expected to occur exactly 2 times, and is always the first activity in that sub-process to be executed. Therefore, the language of the Declare model $X1$ is $\mathcal{L}(X1) = \{\langle s, s \rangle\}$. Model $X2$ is a bit more complex. It contains an $absence2(x)$ constraint for each activity $x \in \{p, q, r, t\}$, but no existence constraints, meaning that all activities can occur either once or not at all. In addition, there is a $chain\ precedence(q, r)$ constraint indicating that r can only occur immediately after q and a $chain\ response(r, t)$ constraint indicating that after r , t must happen immediately. Note that “immediately” is interpreted in the context of the Declare model, so while neither p or q may happen between them, other activities in the hybrid model that can happen in parallel with $X2$ may still occur between them. This means that the language of the Declare model $X2$ is: $\mathcal{L}(X2) = \{\langle \rangle, \langle t \rangle, \langle q \rangle, \langle q, t \rangle, \langle t, q \rangle, \langle q, r, t \rangle, \langle p \rangle, \langle p, t \rangle, \langle t, p \rangle, \langle q, p \rangle, \langle p, q \rangle, \langle p, q, t \rangle, \langle p, t, q \rangle, \langle q, p, t \rangle, \langle q, t, p \rangle, \langle t, p, q \rangle, \langle t, q, p \rangle, \langle p, q, r, t \rangle, \langle q, r, t, p \rangle\}$.

Now that we have introduced its component notations, we can give the language of the entire hybrid model in Fig. 1. Because of the many possible interleavings, giving the full set of all possible traces would be fairly cumbersome, so instead we provide it by using the operators used in Definition 1: $\{\langle a, b \rangle\} \cdot (\mathcal{L}(X1) \diamond (\mathcal{L}(X2) \cdot \{\langle \rangle, \langle o \rangle\})) \cdot \{\langle c, d, e \rangle\}$.

4 Hybrid Model Discovery

Our approach is inspired by two observations: (1) imperative miners in general, and the Inductive miner in particular, have difficulties in dealing with unstructured behavior without sacrificing precision (while maintaining a high fitness value), and (2) declarative models are aimed at describing unstructured behavior in a concise way. This led us to an approach where we first mine an imperative model (a process tree). Afterwards, we try to substitute various sub-processes with declarative models to improve precision without sacrificing fitness. Figure 2

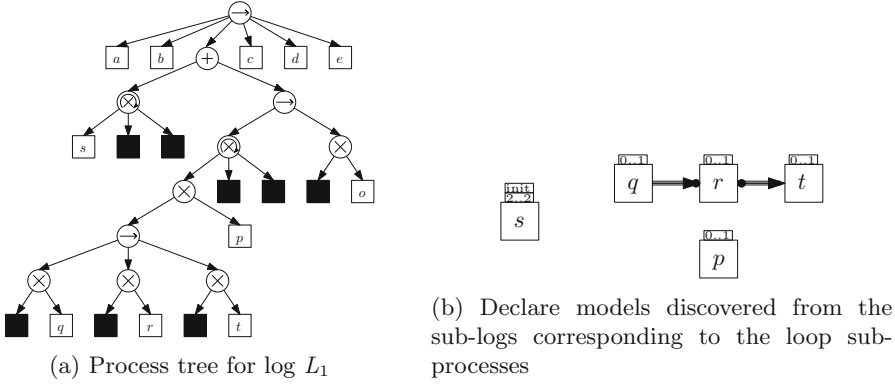


Fig. 3. Intermediate outcomes of the proposed approach

gives a broad overview of the approach, consisting of 4 steps. To exemplify our approach, we use the example log L_1 .

Definition 4 (L_1 : Running Example). Recall that the multiset $[a^4, b^6]$ is a set containing four times element a and six times element b . $L_1 = [\langle a, b, s, s, c, d, e \rangle^5, \langle a, b, o, s, s, c, d, e \rangle^5, \langle a, b, s, o, s, c, d, e \rangle^5, \langle a, b, q, s, r, s, t, c, d, e \rangle^5, \langle a, b, s, s, q, p, c, d, e \rangle^5, \langle a, b, s, q, s, r, t, c, d, e \rangle^5, \langle a, b, s, t, p, s, c, d, e \rangle^5, \langle a, b, s, s, t, q, c, d, e \rangle^5, \langle a, b, p, t, s, q, s, c, d, e \rangle^5, \langle a, b, s, s, p, t, c, d, e \rangle^5, \langle a, b, p, s, q, s, r, t, c, d, e \rangle^5, \langle a, b, s, s, q, t, o, c, d, e \rangle^5]$.

4.1 Step 1: Mine Process Tree

For the first step of our approach, we use the Inductive miner [16] to generate a process tree from the input log. The Inductive miner builds a directly-follows-graph for the unique activities of the log and uses it to find structures and construct a process tree. If the Inductive miner cannot find any structure, it will return a flower model where any behavior is allowed. This makes the Inductive miner ideal for our approach: not only is it one of the most popular mining algorithms, but also it avoids finding structures when none exists and instead simply returns a pocket-of-flexibility, which can then be mined declaratively. We use the standard Inductive miner with a noise threshold of 0. In this way, the generated model is guaranteed to have 100% fitness and we can make a fair comparison between a pure imperative and a hybrid approach, where we only consider improvements on the precision of the model. For our example log, the Inductive miner creates the process tree shown in Fig. 3a.

4.2 Step 2: Select Unstructured Nodes

Once we have mined the process tree, we determine which parts should be replaced with a declarative sub-process. For every permutation of nodes, we

replace those nodes with Declare sub-processes, determine the precision of the model, and return the most precise permutation. If a node is selected to be replaced with a declarative sub-process and also a higher node in the tree is selected, then we only substitute the higher node in the tree, i.e., the node and all its descendants are replaced.

Note that, when we apply our approach to the example log and process tree from Fig. 3a by replacing the loop sub-processes with declarative sub-processes, this results in an increase of precision. This is because the loops allow for any number of occurrences of s and p, q, r, t , while, in reality, the allowed number of occurrences is bound. This bound can be found by the Declare miner.

4.3 Step 3: Mine Declarative Models

For each permutation of nodes to be replaced, we extract the sub-logs corresponding to these nodes. Afterwards, we run a declarative miner on each sub-log to generate a declarative model corresponding to the selected node and its descendants. In our particular case, we use the Declare miner [17] for this task, but we have also made a wrapper for MINERful [13]. The sub-log corresponding to the first loop sub-process is $[\langle s, s \rangle^{60}]$. Passing this log to the Declare miner generates the Declare model in Fig. 3b on the left. The sub-log corresponding to the second loop sub-process is $[\langle \rangle^5, \langle o \rangle^{10}, \langle q, r, t \rangle^{10}, \langle q, p \rangle^5, \langle t, p \rangle^5, \langle t, q \rangle^5, \langle p, t, q \rangle^5, \langle p, t \rangle^5, \langle p, q, r, t \rangle^5, \langle q, t, o \rangle^5]$. Passing this log to the Declare miner generates the Declare model in Fig. 3b on the right.

4.4 Step 4: Construct Hybrid Model

Once we have a Declare model corresponding to each selected node, we replace these nodes and their descendants with an abstract activity representing a sub-process. The sub-process is then defined by the declarative model mined from the sub-log for that node. Applying this technique to the process tree from Fig. 3a and the Declare models from Fig. 3b, we can derive the hybrid model in Fig. 1.

5 Evaluation

Our approach has been implemented in the HYBRID package in ProM. The main plug-in for mining is called “Mine Hybrid Model”. Figure 4 shows the visualization of a hybrid model in the plug-in.

We evaluated our approach on both synthetic and real-life logs. We use the synthetic examples to show that the algorithm is able to detect and construct hybrid models when these are clearly present in the log. We use the real-life logs to demonstrate how well the approach works on actual logs coming from industrial settings, and to discover if any of these have a clear hybrid structure. We have evaluated the output of the Hybrid miner (HM) in terms of precision. This in order to show that a hybrid model indeed can help in making the unstructured parts of the process tree resulting from the Inductive miner (IM) more precise. To show this increase in precision, we introduce the *relative increase in precision* (RIIP) as the percentage of increase in precision.

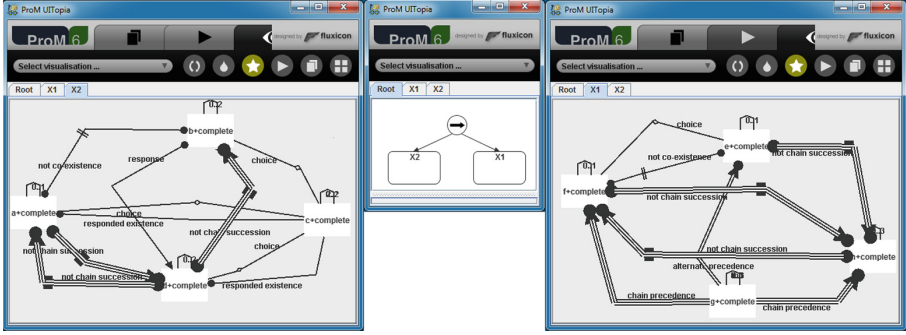


Fig. 4. Hybrid model for log L_3

5.1 Precision

We use a metric based on *escaping edges* [3] for the computation of precision. Our main consideration for choosing this metric is that the precision is defined based on a transition system, a formalism that both a process tree and a Declare model can be transformed into. For every visited state in the transition system, the outgoing transitions used by the log are compared to the total number of possible outgoing transitions. Transitions that are never visited by the log, the *escaping edges*, are counted against the precision of the model. The metric takes into account the frequencies of states, i.e., a frequently visited state has a higher impact on the precision than an infrequently visited state, and states that are never visited by the log are not taken into consideration.

5.2 Synthetic Logs

In order to show the possible gains of using a hybrid approach over solely using the Inductive miner, we have created four different synthetic logs, each serving as a representative for an interesting class of hybrid behavior.

L_1 : *Running example*. The first log we examined was that of the running example, which is an example of a log where part of the behavior is very structured (events a , b , c , d and e), and part is unstructured (events o , p , q , r , s and t) that always occurs in a specific location (between events b and c). The most precise model for L_1 is more complex than the simplified Hybrid model used for illustrative purposes in Fig. 1. By mining the log with the Hybrid miner, we achieve a precision improvement of +0.194 with respect to the Inductive miner. The relative improvement (how much the precision of the hybrid model improved when compared to the precision of the imperative model) RIiP is of 29.59%. The resulting model is clearly hybrid, consisting of a process tree with 7 nodes and a declarative sub-process of 30 constraints.

L_2 : *Running example unbalanced*. In the running example, the traces of the log are balanced (they occur equally frequently). In the second example, we

experimented with making the traces exhibiting unstructured behavior more frequently, with the aim of discovering if this increased the relative precision of the hybrid model. Definition 5 shows the log used.

Definition 5 (L_2 : Running Example Unbalanced). $L_2 = [\langle a, b, s, s, c, d, e \rangle^1, \langle a, b, o, s, s, c, d, e \rangle^1, \langle a, b, s, o, s, c, d, e \rangle^1, \langle a, b, q, s, r, s, t, c, d, e \rangle^{100}, \langle a, b, s, s, q, p, c, d, e \rangle^1, \langle a, b, s, q, s, r, t, c, d, e \rangle^1, \langle a, b, s, t, p, s, c, d, e \rangle^1, \langle a, b, s, s, t, q, c, d, e \rangle^1, \langle a, b, p, t, s, q, s, c, d, e \rangle^{100}, \langle a, b, s, s, p, t, c, d, e \rangle^5, \langle a, b, p, s, q, s, r, t, c, d, e \rangle^{100}, \langle a, b, s, s, q, t, o, c, d, e \rangle^5]$.

For this log, the Hybrid miner provides an improvement from 0.543 to 0.827 with respect to the Inductive miner. Compared to L_1 , the hybrid model accounts for the unstructured parts of the log being more pronounced by adding 2 additional constraints to the declarative sub-process. Note that, while there is an absolute improvement that is slightly better (+0.284) than the one obtained for L_1 , the relative improvement RiP is much higher (52.18%).

L_3 : *Abstract structure.* For the third experiment, we used log L_3 .

Definition 6 (L_3 : Abstract Structure). $L_3 = [\langle a, c, c, d, g, h, e, g, g \rangle^5, \langle c, b, c, d, g, g, f, g, h \rangle^5, \langle c, a, c, d, d, g, h, g, h, g, h, e \rangle^5, \langle d, c, c, a, g, h, g, f, g \rangle^5, \langle c, a, c, g, g, f, g, h \rangle^5, \langle c, c, d, g, h, g, e, g, h \rangle^5, \langle c, b, c, d, g, g, g, f \rangle^5, \langle d, c, b, c, d, g, g, g, e \rangle^5, \langle b, c, c, b, d, g, f, g, g \rangle^5, \langle c, c, g, g, g, e \rangle^5]$.

The log exhibits an abstract structure, i.e., while all the activities of the log occur in an unstructured manner, they can be separated in two sets $\{a, b, c, d\}$ and $\{e, f, g, h\}$, where events of the second set always occur after the events of the first set. As expected the Hybrid miner finds a hybrid model consisting of a sequence flow and two declarative sub-processes, one containing 14 constraints, the other 15. The hybrid model results in a precision improvement from 0.459 to 0.616.

L_4 : *Unstructured sequences.* For the fourth experiment, we used log L_4 .

Definition 7 (L_4 : Unstructured Sequences). $L_4 = [\langle a, b, c, f, g \rangle^5, \langle f, g, a, b, c, d, e \rangle^5, \langle a, b, c, h, i, d, e \rangle^5, \langle a, b, c, f, g, d, e \rangle^5, \langle h, i, a, b, c, h, i, h, i \rangle^5, \langle h, i, a, b, c, d, e, a, b, c \rangle^5, \langle a, b, c, a, b, c, f, g, d, e \rangle^5, \langle f, g, a, b, c, a, b, c, d, e \rangle^5, \langle h, i, a, b, c \rangle^5, \langle h, i, h, i, a, b, c, d, e \rangle^5]$.

This log contains four strict sequences, $\langle a, b, c \rangle$, $\langle d, e \rangle$, $\langle f, g \rangle$ and $\langle h, i \rangle$, which themselves occur in an unstructured manner. This example can be seen as being in the inverse class as L_3 , i.e., while all activities occur in a structured manner, there is a very clear unstructured abstract behavior in the log. Since the best way of representing this process would be by using a declarative model with imperative sub-processes, mining the log with the Hybrid miner (which is developed to build imperative models with declarative sub-processes) results in a fully declarative model. The declarative model provides a precision improvement from 0.743 to 0.897 with respect to the Inductive miner.

Table 2. Evaluation results of the Hybrid miner

	Precision			Hybrid miner		
	IM	HM	RiP	#nodes	#sub-processes	#constraints
L_1	0.6552	0.8491	29.59%	7	1	30
L_2	0.5432	0.8266	52.18%	7	1	30
L_3	0.4595	0.6165	34.17%	3	2	14,15
L_4	0.7428	0.8966	20.70%	1	1	101
BPIC 2012	0.2401	0.6068	152.74%	1	1	1286
BPIC 2013	0.4044	0.4044	0.00%	44	0	0
BPIC 2017	0.0601	0.0601	0.00%	72	1	1
WABO2	0.0049	0.1673	3347%	1	1	78 100
WABO3	0.0059	0.2101	3471%	1	1	75 893
Sepsis cases	0.1905	0.3762	97.47%	1	1	211
RTFM	0.8181	0.9625	17.65%	1	1	73

5.3 Real-Life Logs

We also experimented on a number of real-life logs, in particular the BPI Challenge (BPIC) 2012, 2013 (incidents) and 2017, WABO 2 and 3, Sepsis cases and Road Traffic Fine Management (RTFM). Each of these logs was retrieved from the 4TU repository.¹ As mentioned, in order to find the most precise model, the Hybrid miner tries all permutations of nodes resulting in models ranging from fully imperative via hybridity to fully declarative. In most cases, a fully declarative model was returned as being the most precise. The only exception are the BPIC 2013 and 2017 logs. The latter returns a hybrid model that contains only one Declare constraint that gives no noticeable precision improvement. For the WABO logs, the models found by the Hybrid miner show a precision improvement of over 30 times even if these models are extremely large, consisting of over 75 000 constraints. The BPIC 2012 and Sepsis cases logs also show a significant increase in precision and return more reasonably sized declarative models.

5.4 Discussion

Table 2 shows the detailed results of all the experiments we ran. Overall the Hybrid miner performed quite well on both synthetic and real-life examples. However, in the case of the real-life logs this was mostly due to the fact that the Hybrid miner was able to determine that a fully declarative model would have the highest precision and provided such a model. Granted, this increase in precision came at the cost of simplicity since more than 75 000 constraints cannot be comprehended by an end-user. This is also one of the limitations of

¹ <https://data.4tu.nl/repository/>.

our current approach: getting the right settings on the used miners to give a precise *and* comprehensible result.

In future work, we would like to investigate why the Hybrid miner did not generate any particularly interesting hybrid models from real-life logs. One particularly promising hypothesis is that these logs fit into the category represented by L_4 , which would be better represented by a hybrid model where the top-level is declarative, and sub-processes are imperative. Another explanation could be that the structure of the process tree returned by the Inductive miner hinders further improvements. After all, we take the structure of the process tree as-is without considering other structures that are language equivalent to the returned process tree.

In addition, it should be noted that, while we are using one of the most widely accepted precision metric, there is still an ongoing search for better metrics [33]. For example, for a log containing at most 2 occurrences of an activity, the current metric gives only a small benefit to an *absence3* constraint over a process tree loop, even though, from a language inclusion perspective, the *absence3* constraint is infinitely more precise. Improvements to the way precision is calculated for hybrid models may also lead to different results.

6 Conclusion

We presented a novel technique for mining hybrid models, which combine the strengths of the imperative and declarative process modeling paradigms. We implemented our technique as a ProM plug-in and evaluated the approach on several synthetic and real-life logs. Compared to the Inductive miner, the miner showed significant improvements in precision for both synthetic and real-life logs. In the case of real-life logs, it mostly found that purely declarative models were the most precise, whereas in the case of synthetic logs, proper hybrid models were found. Precision improvements ranged up to 52.18% for synthetic logs and up to 3471% for real-life logs.

The presented approach is easily extendible and customizable. For example, we can already plug any other miner that provides a process tree as output into our implementation and it is also possible to extend the general approach to other types of imperative models and miners. It is also possible to apply different metrics for precision and to use various declarative miners.

As future work, we would like to implement and experiment with different variants of the approach (e.g., different miners and ways to partition behavior in logs into declarative and imperative parts). Moreover, we also plan to use an approach opposite to the one proposed here: we could first attempt to mine a log declaratively, analyze the resulting model to detect parts which are overly constrained, and replace these with imperative sub-processes.

References

1. van der Aalst, W.M.P., Adams, M., ter Hofstede, A.H.M., Pesic, M., Schonenberg, H.: Flexibility as a service. In: Chen, L., Liu, C., Liu, Q., Deng, K. (eds.) DASFAA 2009. LNCS, vol. 5667, pp. 319–333. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-04205-8_27
2. van der Aalst, W.M.P.: Process Mining - Data Science in Action, 2nd edn. Springer, Heidelberg (2016). <https://doi.org/10.1007/978-3-662-49851-4>
3. van der Aalst, W.M.P., Adriansyah, A., van Dongen, B.F.: Replaying history on process models for conformance checking and performance analysis. Wiley Interdisc. Rev. Data Min. Knowl. Disc. **2**(2), 182–192 (2012)
4. van der Aalst, W.M.P., Rubin, V.A., Verbeek, H.M.W., van Dongen, B.F., Kindler, E., Günther, C.W.: Process mining: a two-step approach to balance between underfitting and overfitting. Softw. Syst. Model. **9**(1), 87–111 (2010)
5. van der Aalst, W.M.P., Weijters, A.J.M.M., Maruster, L.: Workflow mining: discovering process models from event logs. IEEE Trans. Knowl. Data Eng. **16**(9), 1128–1142 (2004)
6. Alberti, M., Chesani, F., Gavanelli, M., Lamma, E., Mello, P., Torroni, P.: Verifiable agent interaction in abductive logic programming: the SCIFF framework. ACM Trans. Comput. Log. **9**(4), 29:1–29:43 (2008)
7. Back, C.O., Debois, S., Slaats, T.: Towards an entropy-based analysis of log variability. In: Teniente, E., Weidlich, M. (eds.) BPM 2017. LNBIP, vol. 308, pp. 53–70. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-74030-0_4
8. Chesani, F., Lamma, E., Mello, P., Montali, M., Riguzzi, F., Storari, S.: Exploiting inductive logic programming techniques for declarative process mining. T. Petri Nets Other Models Concurrency **2**, 278–295 (2009)
9. De Giacomo, G., Dumas, M., Maggi, F.M., Montali, M.: Declarative process modeling in BPMN. In: Zdravkovic, J., Kirikova, M., Johannesson, P. (eds.) CAiSE 2015. LNCS, vol. 9097, pp. 84–100. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-19069-3_6
10. Debois, S., Hildebrandt, T., Slaats, T.: Hierarchical declarative modelling with refinement and sub-processes. In: Sadiq, S., Soffer, P., Völzer, H. (eds.) BPM 2014. LNCS, vol. 8659, pp. 18–33. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-10172-9_2
11. Debois, S., Hildebrandt, T.T., Marquard, M., Slaats, T.: Hybrid process technologies in the financial sector. In: BPM (Industry track), pp. 107–119 (2015)
12. Debois, S., Hildebrandt, T.T., Slaats, T., Marquard, M.: A case for declarative process modelling: Agile development of a grant application system. In: EDOC Workshops, vol. 14, pp. 126–133 (2014)
13. Di Ciccio, C., Mecella, M.: A two-step fast algorithm for the automated discovery of declarative workflows. In: CIDM, pp. 135–142. IEEE (2013)
14. Lamma, E., Mello, P., Montali, M., Riguzzi, F., Storari, S.: Inducing declarative logic-based models from labeled traces. In: Alonso, G., Dadam, P., Rosemann, M. (eds.) BPM 2007. LNCS, vol. 4714, pp. 344–359. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-75183-0_25
15. Lamma, E., Mello, P., Riguzzi, F., Storari, S.: Applying inductive logic programming to process mining. In: Blockeel, H., Ramon, J., Shavlik, J., Tadepalli, P. (eds.) ILP 2007. LNCS (LNAI), vol. 4894, pp. 132–146. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-78469-2_16

16. Leemans, S.J.J., Fahland, D., van der Aalst, W.M.P.: Discovering block-structured process models from event logs - a constructive approach. In: Colom, J.-M., Desel, J. (eds.) *PETRI NETS 2013*. LNCS, vol. 7927, pp. 311–329. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-38697-8_17
17. Maggi, F.M.: Declarative process mining with the declare component of ProM. In: *BPM (Demos)*. CEUR Workshop Proceedings, vol. 1021. CEUR-WS.org (2013)
18. Maggi, F.M., Bose, R.P.J.C., van der Aalst, W.M.P.: Efficient discovery of understandable declarative process models from event logs. In: Ralyté, J., Franch, X., Brinkkemper, S., Wrycza, S. (eds.) *CAiSE 2012*. LNCS, vol. 7328, pp. 270–285. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-31095-9_18
19. Maggi, F.M., Mooij, A.J., van der Aalst, W.M.P.: User-guided discovery of declarative process models. In: *CIDM*, pp. 192–199. IEEE (2011)
20. Maggi, F.M., Slaats, T., Reijers, H.A.: The automated discovery of hybrid processes. In: Sadiq, S., Soffer, P., Völzer, H. (eds.) *BPM 2014*. LNCS, vol. 8659, pp. 392–399. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-10172-9_27
21. Marquard, M., Shahzad, M., Slaats, T.: Web-based modelling and collaborative simulation of declarative processes. In: Motahari-Nezhad, H.R., Recker, J., Weidlich, M. (eds.) *BPM 2015*. LNCS, vol. 9253, pp. 209–225. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-23063-4_15
22. de Medeiros, A.K.A., Weijters, A.J.M.M., van der Aalst, W.M.P.: Genetic process mining: an experimental evaluation. *Data Min. Knowl. Disc.* **14**(2), 245–304 (2007)
23. Montali, M.: Specification and Verification of Declarative Open Interaction Models. *LNBIP*, vol. 56. Springer, Heidelberg (2010). <https://doi.org/10.1007/978-3-642-14538-4>
24. Pesic, M., Schonenberg, H., van der Aalst, W.M.P.: DECLARE: full support for loosely-structured processes. In: *EDOC*, pp. 287–300 (2007)
25. Pichler, P., Weber, B., Zugal, S., Pinggera, J., Mendling, J., Reijers, H.A.: Imperative versus declarative process modeling languages: an empirical investigation. In: Daniel, F., Barkaoui, K., Dustdar, S. (eds.) *BPM 2011*. LNBIP, vol. 99, pp. 383–394. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-28108-2_37
26. Reijers, H.A., Slaats, T., Stahl, C.: Declarative modeling—an academic dream or the future for BPM? In: Daniel, F., Wang, J., Weber, B. (eds.) *BPM 2013*. LNCS, vol. 8094, pp. 307–322. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40176-3_26
27. Sadiq, S., Sadiq, W., Orłowska, M.: Pockets of flexibility in workflow specification. In: S.Kunii, H., Jajodia, S., Sølvberg, A. (eds.) *ER 2001*. LNCS, vol. 2224, pp. 513–526. Springer, Heidelberg (2001). https://doi.org/10.1007/3-540-45581-7_38
28. Schönig, S., Rogge-Solti, A., Cabanillas, C., Jablonski, S., Mendling, J.: Efficient and customisable declarative process mining with SQL. In: Nurcan, S., Soffer, P., Bajec, M., Eder, J. (eds.) *CAiSE 2016*. LNCS, vol. 9694, pp. 290–305. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-39696-5_18
29. Schunselaar, D.M.M.: Configurable Process Trees: Elicitation, Analysis, and Enactment. Ph.D. thesis, Eindhoven University of Technology (2016)
30. Slaats, T., Schunselaar, D.M.M., Maggi, F.M., Reijers, H.A.: The semantics of hybrid process models. In: *OTM CoopIS*. pp. 531–551 (2016)
31. Smedt, J.D., Weerd, J.D., Vanthienen, J., Poels, G.: Mixed-paradigm process modeling with intertwined state spaces. *Bus. IS Eng.* **58**(1), 19–29 (2016)
32. Smedt, J.D., Weerd, J.D., Vanthienen, J.: Fusion miner: process discovery for mixed-paradigm models. *Decis. Support Syst.* **77**, 123–136 (2015)
33. Tax, N., Lu, X., Sidorova, N., Fahland, D., van der Aalst, W.M.P.: The imprecisions of precision measures in process mining (2017). <https://arxiv.org/abs/1705.03303>

34. Weijters, A.J.M.M., van der Aalst, W.M.P.: Rediscovering workflow models from event-based data using Little Thumb. *Integr. Comput.-Aided Eng.* **10**(2), 151–162 (2003)
35. Westergaard, M., Slaats, T.: CPN Tools 4: A process modeling tool combining declarative and imperative paradigms. In: *BPM (Demos)* (2013)
36. Westergaard, M., Slaats, T.: Mixing paradigms for more comprehensible models. In: Daniel, F., Wang, J., Weber, B. (eds.) *BPM 2013. LNCS*, vol. 8094, pp. 283–290. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40176-3_24



Extending BPSim Based on Workflow Resource Patterns

Nehal Affi, Ahmed Awad^(✉), and Hisham M. Abdelsalam^(✉)

Faculty of Computers and Information, Cairo University, Giza, Egypt
n.affi@grad.fci-cu.edu.eg, {a.gaafar,h.abdelsalam}@fci-cu.edu.eg

Abstract. The improper modeling of human resources spotted within business process simulation approaches affect appropriate evaluation of dynamic business processes behavior over time. The BPSim standard is acknowledged as a first step towards streamlining the experience of business process simulation and providing a tool independent exchange format for so-called simulation scenarios. Unfortunately, BPSim is not fully elaborated regarding the resource perspective. This paper introduces RBPSim, an extension of BPSim inspired by the well-known Workflow Resource Patterns. Moreover, the paper provides refined inter and intra-relationships among Workflow Resource Patterns to guide a modular implementation within simulation tools.

Keywords: BPSim · Workflow resource patterns
Business process simulation · Human resource perspective
Business processes

1 Introduction

Business process simulation (BPS) involves performing simulation experiments through analyzing the dynamic behavior of business processes over time with respect to performance metrics such as cycle time, cost and resources utilization [1, 2]. For simulating business processes, we need to model at least three perspectives [3, 4]: control-flow, data and resources perspective. One of the main limitations affecting the current BPS approaches is the naïve modeling of resources perspective [3]. The term Resource is used hereinafter to specify human resources.

The resource perspective captures the resources representation, distribution and authorization [5] with respect to work preference, speed and realistic allocation elements and extension points [5, 6]. It is important to ensure the proper distribution of suitable resources to carry out work items execution [7] to avoid unsuccessful simulation experiment [8, 9]. Yet it is hard to specify how resources are properly represented and how the work items are distributed among eligible resources.

To accomplish an effective resources representation within the simulation experiment, several requirements [10] should be defined such as: one or more *required resources* to handle work items execution; *resource availability* to execute work items [8, 10]; *execution duration* of work items are not constant and should follow a probabilistic distribution [10]; *resource share-ability*, as resources are not dedicating all their time to one work item and may divide time simultaneously between different work items [8]; *context switching overhead* resources may require time intervals between different work items execution [10]; and resources should have the ability for *work items selection* based on their own characteristics or preferences as well work items should include a defined priority to specify how resources will select the order of work items from their working queue based on FIFO, LIFO or priority [10].

The Business Process Simulation (BPSim) 2.0 standard [11], developed by WfMC¹, allows BPMN [12] models to be augmented with simulation-specific parameters such as tasks durations, branching probabilities, case arrival rates, etc. BPSim meta-model is not fully elaborated regarding the resources perspective [10, 13]. However, it provides an extension mechanism for attaching additional attributes and elements to the original meta-model [11].

In a previous work [14], we suggested a preliminary extension of BPSim to make it resource-aware, RBPSim. The extension is inspired by the well known workflow resource patterns [9]. In this paper we provide a detailed description of RBPSim, our proposed extension of BPSim 2.0 standard presented in [14], to allow sophisticated modeling of resource perspective within the simulation scenarios. First, we contribute a deeper discussion of the Workflow Resource Patterns [9] and provide fine-grained inter and intra-relationships among those patterns. Next, we present how RBPSim meta-model will express resource representation, selection and distribution using the refined complement relationships among Workflow Resource Patterns that shall help modularize the implementation support of such patterns within simulation tools.

The rest of the paper is organized as the following: Sect. 2 revisits the Workflow Resource Patterns and further refines their relationships to modularize their use in simulation or process execution. Section 3 introduces our extension of BPSim with the refined Workflow Resource Patterns. Section 4 discusses the related work and Sect. 5 concludes the paper with an outlook on future work.

2 Revisiting Workflow Resource Patterns

Distinct categories of Workflow Resource Patterns define different requirements to capture resources perspective in business process models [15]. It is helpful to study the relationships among those patterns to build a minimal set of constructs that can be reused in the related patterns. Previous studies [9] about such relationships introduced a simple “related to” relationship between pairs of patterns either across or within the same pattern group. Here, we refined those relationships using *complement-perspective*.

¹ <https://www.wfmc.org/>.

2.1 Creation Patterns Intra Relationships

The intra-relationships among Creation Patterns are summarized in Fig. 1. They are: the *is-a* relationship, the *complements* relationship, and the *opposite to* relationship to indicate that patterns are the opposite or negation of each other, precautions have to take place when those opposite patterns are applied on an overlapping set of tasks. We have three patterns at the top of the hierarchy: Role-based Distribution, Authorization and Retain Familiar. The other Creation Patterns are special cases of those patterns as follows:

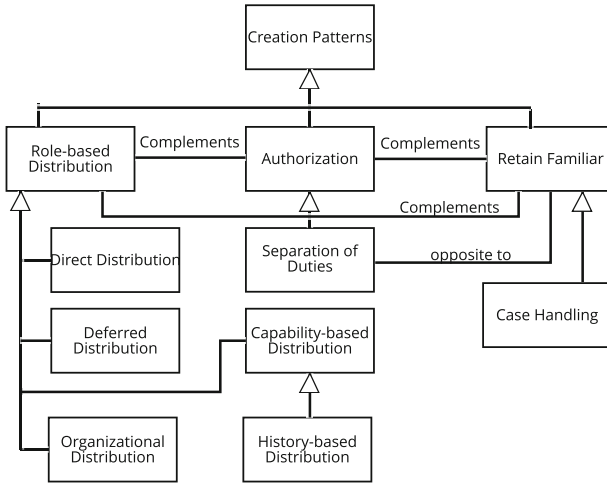


Fig. 1. Creation patterns intra-relationships

Direct Distribution can be modeled as Role-based Distribution by making sure that only one resource is a member of the role. *Deferred Distribution* is a special case of Role-based Distribution through specifying the role at design time but the actual members will be resolved at run (simulation) time. However, the resolution logic at runtime must be specified in the simulation model, this can fall back to other distribution patterns. *Organizational Distribution* is special case of Role-based Distribution because we can simulate the organizational position as a role.

Capability-based Distribution is a special case of Role-based as a virtual role can collect all resources that possess those capabilities. The capability constraints are evaluated against the whole set of resources and matching individuals are added to the virtual role. This step can take place either before running the simulation, in the form of preparing simulation data or can be done as one of the initialization steps of the simulation model. In the latter case, the selection logic is part of the simulation model and is run prior to the very first task in the process. *History-based Distribution* is a special case of Capability-based

Distribution where we can add a property `Number of completed items` along with its value to the set of capabilities to look for. The value shall be supplied as part of the simulation parameters.

Authorization according to [16], Authorization touches both design- and runtime aspects of resource accessibility to a work item. Looking from the design time perspective, Authorization adds a security framework that complements the work item distribution to resources. That is why we put it as complementary to both Role-based Distribution and Retain Familiar patterns as it is possible to be used for refining the allocation. From the runtime perspective, Authorization is concerned with the actions a resource can take with respect to work items in his list like starting, suspending, delegating etc. For simulation purposes, we are concerned with the design time perspective and we model it in the same way we model the Role-based Distribution pattern.

Separation of Duties is considered as a special case of the Authorization pattern. Our argument here is that Separation of Duties states that the performer of task *B* cannot be the same performer of task *A* within the same case. In that sense, at runtime performer of task *A* in a given case is *not authorized* to perform task *B* for the same case. Separation of Duties is also considered *opposite to Retain Familiar pattern* where in Retain Familiar the resource is executing the whole case while in Separation of Duties resource is allowed to execute part of the case depending on some specification. Thus, it is a contradiction to have both patterns applied on the same process model. *Case Handling* is a special case of Retain Familiar where all work items of the case are distributed to the same resource.

2.2 Push Patterns Intra Relationships

Push Patterns consists of nine patterns subdivided into three sets: as in Fig. 2.

Resource Distribution is concerned with the offering of work item to a single or multiple resources with the possibility to accept or reject offered work items. From simulation point of view, acceptance/rejection of offered work items could be established using a predefined probability. *Resource Allocation* identifies the suitable resource when many possible candidates are available and then allocating work item based on Random allocation, Round Robin or Shortest Queue. Resource Allocation complements Distribution by Allocation-Single Resource pattern where work item is allocated to a single resource among convenient candidates based on the resource's queue length (shortest queue), cyclic basis or randomly. *Distribution Time* spots the time of work item distribution to an appropriate resource. The Early Distribution pattern indicates the ability of resources to view future work item and provides "booking a resource"? mechanism for upcoming work items. Both Distribution on Enablement and Late Distribution patterns correspond to distributing enabled work items for execution to a resource. Distribution Time complements Resource Selection as whatever the selected distribution time, work item requires to be offered and allocated to a convenient resource and distinguishing them at simulation time would not make a difference.

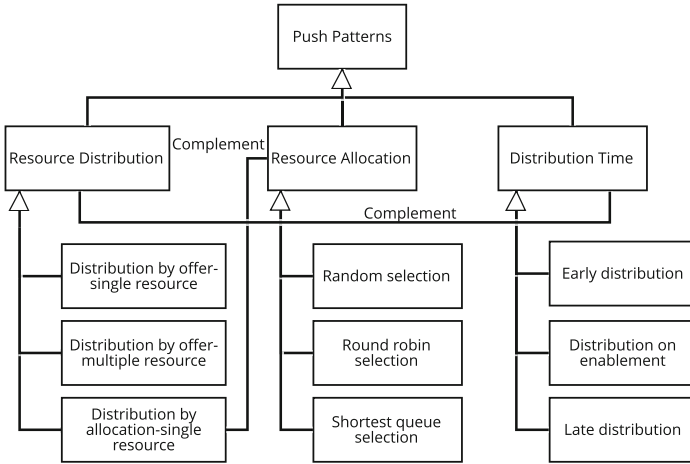


Fig. 2. Push patterns intra-relationships

2.3 Pull Patterns Intra Relationships

Pull Patterns are concerned with resource awareness of work items available for execution either through direct system offering or a shared work list [15]. Pull Patterns contain six patterns as shown in Fig. 3 and are divided into two sets based on *Pull Action* specified by the resource and *Workitems Management* that display work items for the resources and impact the execution sequence either by the system or by the resource. Selection Autonomy complements all the Pull Action transitions where the resource preference or specific characteristics affect the resource’s commencement of offered or allocated work items.

2.4 Workflow Resource Patterns Inter-relationships

Here, we provide a refined inter-relationship among Workflow Resource Patterns groups. Table 1 summarizes complements relationships among Creation and Push Patterns. *Direct Distribution* specifies the identity of the resource that will handle work item execution. That complements the distribution of work item either through offering to a single resource or allocating the work item on a binding basis. In case of accepting offered, work item is allocated to the resource’s work queue. Direct Distribution complements *Early Distribution* requires specifying the resource identity to execute work items in order to distribute enabled work items directly.

Role-based, Organizational, Capability-based and *History-based Distribution* are concerned with either allocating enabled work item to a single resource on a binding basis or offering work items to multiple resources based on their role, organizational position, capabilities, or previous execution history. Among the eligible candidates, the resource allocation would follow cyclic basis, randomly,

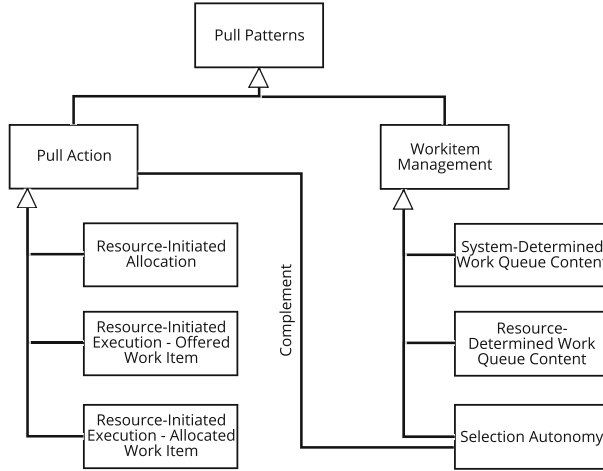


Fig. 3. Pull patterns intra-relationships

Table 1. Complement relationships between Creation and Push Patterns

Creation Patterns	Complement (used with) Push Patterns
Direct Distribution	Distribution by Offer - Single Resource Distribution by Allocation - Single Resource Distribution on Enablement Early Distribution
Role-based Distribution	Distribution by Offer - Multiple Resources
Organizational Distribution	Distribution by Allocation - Single Resource
Capability-based Distribution	Distribution on Enablement
History-based Distribution	Random Allocation Round Robin Allocation Shortest Queue
Deferred Distribution	Late Distribution
Separation of Duties	Distribution by Offer - Single Resource Distribution by Offer - Multiple Resources Distribution by Allocation - Single Resource Distribution on Enablement Random Allocation Round Robin Allocation Shortest Queue
Case Handling	Distribution by Allocation - Single Resource
Retain Familiar	Distribution on Enablement

or shortest queue. *Deferred Distribution* concern with postponing the identification of the required resource for work item execution until run time using *Late Distribution*.

Considering the different types of *Separation of Duties* discussed in [17]. The ability of distributing two work items to distinct resources in a given case is determined through either allocating work item to a single resource r_{alloc} or offering the work item to a single or multiple resource based of specified selection criteria, the accepted work item is allocated to a single resource r_{alloc} based on shortest queue, randomly, or on cyclic basis. The second work item is offered to a set of candidate resources R_{cp} where $r_{alloc} \notin R_{cp}$. Both *Case Handling* and *Retain Familiar* complement *Distribution by Allocation-Single Resource* on binding basis within a given case or a former work item.

Authorization is a Creation Pattern that defines a range of privileges or actions regarding work items execution. It was excluded from Table 1 and explained separately in Table 2 as it complements five groups of resource patterns (Pull, Detour, Auto-Start, Visibility and Multiple Resource).

Table 2. Complement relationships between Authorization and other resource patterns groups

Resource Patterns Group	Complement (used with) Resource Patterns
Pull Patterns	Resource-Initiated Allocation Resource-Initiated Execution - Allocated Work Item Resource-Initiated Execution - Offered Work Item Resource-Determined Work Queue Content Selection Autonomy
Detour Patterns	Delegation, Deallocation, Stateful Reallocation, Skip Stateless Reallocation, Suspension-Resumption and Redo
Auto-start Patterns	Commencement on Creation Commencement on Allocation Piled Execution Chained Execution
Visibility Patterns	Configurable Unallocated Work Item Visibility Configurable Allocated Work Item Visibility
Multiple Resource Patterns	Simultaneous Execution

Authorization complements Pull Patterns among many situations: 1. Commitment of the resource to execute previously allocated work item without the need to immediately work on them; 2. Start working on allocated work items; 3. Immediate execution of allocated work items; 4. Specifying the working queue format; and 5. Executing work items based on resources own preferences.

Interruption of allocated work items occur either by the system or the resource. Complement relationship among Authorization and Detour Patterns

handles a range of possible scenarios including: 1. Delegate started work item to another resource; 2. Deallocate work items that have not been initiated and allow Re-allocation of them; 3. Stateful Reallocate of commenced work items and retain their execution state; 4. Stateless Reallocate of commenced work item to another resource without keeping the current state; 5. Suspend and Resume work items during execution; 6. Skip work items allocated to a resource; 7. Redo execution of previously completed work items.

Authorization complements Auto-Start Patterns through start working on created or allocated work items, initiate next instance of allocated work items by entering the piled execution mode, or automatically start the next work item once the previous one is completed through the chained execution. The ability of a resource to view allocated and unallocated work items could be handled through a complement relationship among Authorization and Visibility Patterns. Authorization complements Multiple Resource Patterns through handling the concurrent execution of more than one work item simultaneously using *Simultaneous Execution*.

3 Extending BPSim with Workflow Resource Patterns

In this section we discuss the level of support offered by BPSim 2.0 standard to resources perspective required for resources representation and distribution. Next, we discuss in detail our extended version of BPSim 2.0 meta-model. The extended meta-model is introduced to refine the spotted limitations in BPSim standard concerning the resources classification and distribution with respect to the refined Workflow Resource Patterns introduced in Sect. 2.

3.1 Business Process Simulation Standard

BPSim 2.0 introduces *Scenario* entity that shall contain all parameters needed to run a simulation experiment [11]. A scenario could be used to capture input parameter specification, simulation results and historical data from past business processes execution [11].

The *ElementParameter* is the concrete class definition of all parameters [11]. It indicates the reference to BPMN elements retrieved from the business process model interface and identifies different elements using a unique identifier. *ElementParameter* is expanded with several parameters [11] such as *TimeParameter*, *ControlParameter* that defines control flow of BPMN element, *CostParameter*, *PriorityParameter* and *ResourceParameter*. Expressions could be attached to *ElementParameter* to specify added functions such as *getResource* to select a collection of available resources, *getResourceByRole* to select a collection of available resources based on role and *Resource* to select an alternative list of available resources.

ResourceParameter [11], groups all parameters related to resources required by business process model element, e.g. task. The parameter determines

resource's availability using *availability* Boolean attribute according to a pre-defined calendar. The default value is True. The required resource quantity to execute a work item is determined using *quantity.selection*, an override of the BPMN ResourceRole element, defines the resource selection criteria based on a defined resource role list or resourceID referring to a specific resource. Among the resource perespective requirements, only determining specific resource, resource role, resource availability and work items execution duration are handled by BPSim 2.0.

TimeParameter and ControlParameter are not valid for resources and resourcesRole elements [11]. However, in limited situations concerning results requests, TimeParameter could be added to resource or resourceRole BPMN elements [11]. CostParameter could be applied to ResourceElement specifying the cost of resources either to fixedCost or unitCost [10]. PriorityParameter and PropertyParameter are not applicable for resource or resourceRole elements [13]. Condition in ControlParameter is just a Boolean attribute and only one condition could be applied to BPMN element. A modification is required to handle specifying one or more conditions required to filter the resource required by the BPMN element, e.g. task [11].

3.2 Resources Definition and Distribution in RBPSim

The extended meta-model "RBPSim" is explained in Fig. 4 and the new added classes are highlighted in grey. RBPSim extension outlines two aspects to define the resources perspective: The *Resources Definition* specifying resources classification and associated privileges and the *Resources Distribution* constraints.

The Resource Definition aspect defines the following classes to specify the resources classification and authorization within the scenario that references to a single business process model. **Resources**, the parent for both **Human** and **Non-Human** Resources, identify the resources classification parameters including a *unique identifier* for each resource, resource *name*, resource *cost* either per use or per defined time interval (if resource cost is specified then cost of ElementParameter need to be adjusted) and the *lagTime* that follows a probabilistic distribution to define the resources pause time interval between work items execution. **Non-Human Resources** enumerates the Resource with two attributes: *resourceType* and *consumptionRate*. *capabilities* is used in **HumanResources** to define the resource's capabilities e.g. years of experience, skills, etc. *canDelegate* is a Boolean attribute used to identify if the resource can delegate work items to another eligible resource(s). **ResourceQueue** used to organize work items waiting for execution. ResourceQueue specify the *maximumQueueLength*, the *queueSortPerformer* that identify the authorized entity for the queue sorting preference (either the system or the resource) and the *sortMethod* to specify the queue sorting method from FIFO, LIFO, or work item priority.

Human resource is a member of **Role** which generalizes over **OrganizationalGroup** and **Position**. Each *Role* may one or more **ResourcePrivileges** that are specified using *ResourcesPrivileges* enumeration and setting the *allowPrivileges* to true. **Resources** may have previous execution history of work

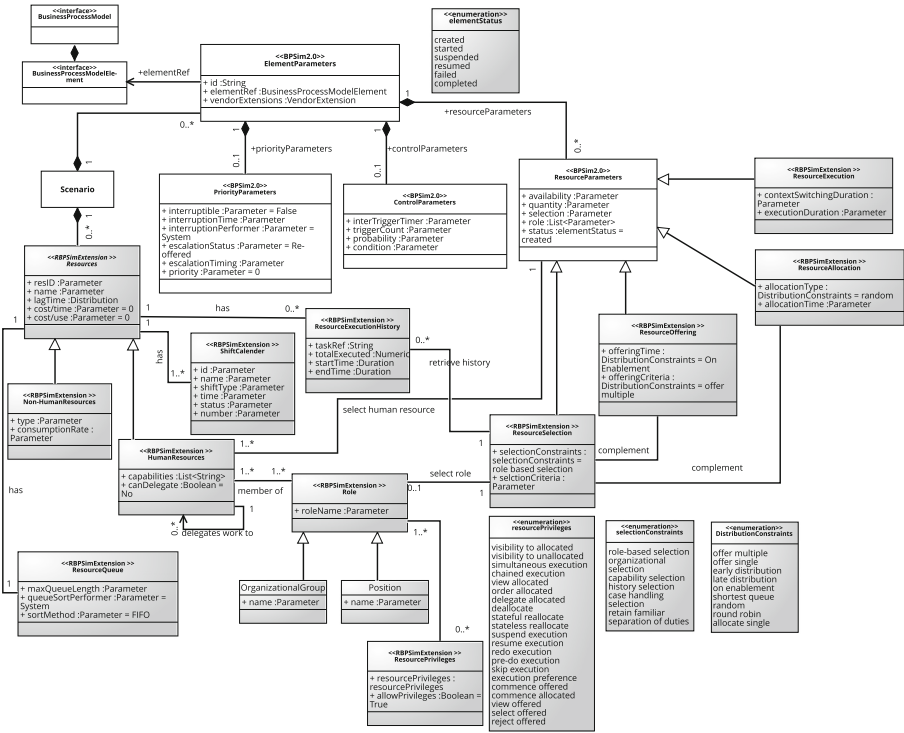


Fig. 4. Extended RBPSim meta-model based on workflow resource patterns

items defined through **ResourceExecutionHistory** and a **ShiftCalendar** indicating the resource availability. **ShiftCalendar** extends BPSim 2.0 Calendar Parameter [11].

The Resources Distribution aspect defines the four phases executing distributed work items. The specified phases are: the *resource selection* to restrict the range of resources that can execute work item, the *resource offering* to determine how and when work items are offered to resources, the *resource allocation* to specify allocation constraints, and the *resource execution*.

ResourceSelection extends the *selection* attribute in **ResourceParameters** and determines the required selection criteria through *selectionConstraints* and *selectionCriteria*. The selection may occur based on resource role, organizational unit, resource capabilities, execution history or handled cases. **ResourceOffering** defines two attributes: *offeringTime* to specify the timing of task offering on task enablement, early offering or late offering, and the *offeringCriteria* to determine either work items are offered to single or multiple resources. **ResourceAllocation** defines the allocation constraints. Allocation follows random allocation, round robin, and resources with shortest queue. **ResourceExecution** determines the *contextSwitchingDuration* a time interval required by the resource between work items execution and the *executionDuration* for work items.

PriorityParameters specify the following priority attributes: *interruptible* Boolean attribute to specify whether the task execution is interruptible, the default value is false. Two new added attributes are extending the interruption specifications with *interruptionTime* and *interruptionPerformer*. The default performer is the system and resource could also be assigned as an interrupter. *escalationStatus* attribute determines the status of escalation to either offered or allocated. *escalationTiming* determines the time of work item escalation.

3.3 Realizing Workflow Resources Patterns in RBPSim

This subsection introduces how the refined inter- and intra-relationships among Workflow Resource Patterns, introduced in Sect. 2, are supported by RBPSim, the new extended BPSim meta-model.

Creation Patterns: *Role-based*, *Capability-based*, *History-based* and *Organizational Distribution* are determined by specifying the selectionConstraints and selectionCriteria in ResourceSelection to eligible resources having specific role, capabilities, history or organizational unit respectively through the connection between ResourceSelection and Role, ResourceExecutionHistory and HumanResources.

Direct Distribution is modeled as Role-based Distribution by making sure that only one resource is a member of the role within Role. Direct Distribution Complements Push patterns through setting offeringCriteria to “offer single” and allocationType to “allocate single” respectively. *Deferred Distribution* is determined by setting the allocationTiming to “early Distribution?”. *Separation of Duties*, *Retain Familiar* and *Case Handling* are determined through defining the required *selectionConstraints* to the required selection pattern and use the execution history to select the appropriate resource.

Push Patterns: The complement relationships between Creation and Push patterns, shown in Table 1, define the offering constraints through specifying the offeringCriteria in ResourceOffering to offer single or multiple resources and the offeringTiming to distribution on enablement. ResourceAllocation defines work items allocation on cyclic basis, randomly, or shortest queue using the allocationType.

Authorization specifies the resources privileges through resourcePrivileges that enumerates the required authorization to handle work items execution. Authorization complements **Pull patterns** through granting the following ResourcePrivileges to resource: view offered, reject offered, view allocated, execution preference and order allocated. **Detour patterns** complements Authorization through identifying delegate, reallocate stateful, reallocate stateless, suspend, resume, skip and redo. Authorization complements **Auto-Start Patterns** by defining commence created, commence allocated, pilled and chained execution. **Visibility Patterns** are determined through visibility to allocated and unallocated. **Multiple Resource Patterns** complements Authorization through simultaneous execution.

3.4 Example

The example explains a simple business process for “Car Maintenance”, see Fig. 5. When cars arrive, an administration employee receives and records car information, the selection, offering and allocation of resource is following those patterns: resources selection requires *Role-based Distribution*, offering follow *Distribution by Offer-Multiple Resources*, distribution timing is based on *Distribution on Enablement* and allocation is based on *Round Robin*. The car is then sent to the mechanical department where an engineer is selected based on capabilities with experience of 3 years and possession of a certificate following *Role-based Distribution* and *Capability-based Distribution*. Finally, an accountant receives and records payments for the work done. The accountant pulls work items from his working queue based on work item priority.




Fig. 5. Car maintenance process in BPMN

Listing 1 is an excerpt of the XML for the simulation scenario parameters based on the extended meta-model from Sect. 3 that realizes the process and patterns discussed above. ElementParameter of “Receive & record car information” task is defined in lines 4 to 17 including the new updates of PriorityParameters in line 6 – 7. ResourceParameter for “Receive & record car information” element is defined in lines 8 to 16 specifying creation pattern (role-based distribution), push patterns (resource offering, resource allocation and distribution timing). ElementParameter for “Fix mechanical issues” task is defined in lines 18 to 30. Resources defining a sample resources participating in the scenario are defined in lines 37 to 50 specifying resource id, name, capabilities, lagTime, resource role, privileges and resource queue properties.

```

1 < bpsim: Scenario id="S1" name="Scenario1:Car Maintenance Model" ...>
2 < bpsim: ScenarioParameters baseTimeUnit="min" />
3 ...
4 < bpsim: ElementParameters elementRef="Receives & records car information">
5 ...
6 < bpsim: PriorityParameters interruptible="True" interruptionTiming="PT15M"
7   escalationStatus="Re-offer" />
8 < bpsim: ResourceParameters elementStatus="created" quantity=1>
9 < bpsim: ResourceSelection selectionConstraints="role-based selection" selectionCriteria="
10   Administration Employee" />
11 < bpsim: ResourceOffering Timing="on enablement" offeringType="offer multiple" />
12 < bpsim: ResourceAllocation allocationType="round robin" />
13 < bpsim: ResourceExecution>
14 < bpsim: contextSwitchingDuration>
15 < UniformDistribution min="3" max="10" />
16 < /bpsim: contextSwitchingDuration>
17 < /bpsim: ResourceParameters>
18 < /bpsim: ElementParameters>
19 < bpsim: ElementParameters elementRef="Fix mechanical issues" >
20 ...
  
```

```

20 < bpsim:ResourceParameters role="Maintenance Engineer" quantity=1>
21 < bpsim:ResourceSelection selectionConstraints="capability selection">
22 < bpsim:selectionCriteria/>
23 < bpsim:Capability experience="3years" certificate="yes"/>
24 </bpsim:selectionCriteria>
25 </bpsim:ResourceSelection>
26 < bpsim:ResourceOffering Timing="on enablement" offeringType="offer multiple"/>
27 < bpsim:ResourceAllocation allocationType="shortest queue"/>
28 < bpsim:ResourceExecution/>
29 </bpsim:ResourceParameters>
30 </bpsim:ElementParameters>
31 < bpsim:ElementParameters elementRef="Receive & record payments">
32 ...
33 < bpsim:ResourceParameters role="Accountant" quantity=1>
34 ...
35 </bpsim:ResourceParameters>
36 </bpsim:ElementParameters>
37 < bpsim:Resources>
38 < bpsim:HumanResources id="Res1" name="Adam" costUnit="$20/h" canDelegate="No"
39 >
40 < bpsim:LagTime>
41 < bpsim:UniformDistribution min="5" max="10">
42 </bpsim:LagTime>
43 < bpsim:Capabilities experience="3 years" certificate="yes"/>
44 < bpsim:Role roleName="Maintenance Engineer"/>
45 < bpsim:ResourceQueue maxQueueLength=10 queueSortPerformer="Resource" sortType
46 ="priority"/>
47 < bpsim:ResourcePrivileges resourcePrivileges="chained execution"
48 allowPrivileges="True"/>
49 < bpsim:ShiftCalender/>
50 < bpsim:ResourceExecutionHistory/>
51 </bpsim:HumanResources>
52 </bpsim:Resources>
53 ...
54 </bpsim:Scenario>

```

Listing 1. RBPSim XML for the Car maintenance model

4 Related Work

Several approaches, e.g. [18–20], discuss the need for business process simulation tools as integral part of business process management.

BPSim was introduced in [11] as standardization effort of process simulation scenario definition. Detailed analysis of BPSim limitations and chances was discussed in [13]. One of the critical limitations mentioned was BPSim resource model elaboration for simulation purposes that require extending BPSim meta-model regarding resources perspective. RAL (Resource Assignment Language) was introduced in [21], a metamodel that could be used with BPMN 2.0 for better resources assignment. Discussions in [13] and in [21] declared that BPSim is partially supported by current tools and requires additional extensions regarding the vendors. Both researches indicated that neither BPSim nor BPMN are fully supporting resources perspective. Work introduced in [5] specifies the resources perspective in BPMN 2.0 metamodel and presents an extension to support the resource modeling with respect to Workflow Resource Patterns. The extension was extended in [22] to provide a conceptual modeling of the BPMN extension using XML that can be processed by any BPMN tools.

A detailed metamodel for resource allocation constraint was introduced in [23]. The introduced metamodel could be reflected in BPMN to overcome drawbacks of supporting resources allocation. Authors in [24] discuss a standard that might be extended to a process analytics framework. The proposed standard would limit the challenges when using simulation for business processes. [3, 25] discuss the limitations of business process simulation usage in-reality and

categorize the risks of current simulation settings. These pitfalls were further extended and refined in [26]. Amongst such risks is the inadequate modeling of resource behavior and availability. Our work is contributing to the avoidance of this risk by means of pattern-based enforcement of resource behavior constraints. Moreover, in [25], the authors use CPN tools by example to address the pitfalls identified where they only addressed modeling of resource availability without constraining resource allocation as we discuss in this paper.

A simulation environment based on YAWL workflow tool and CPN tools was discussed in [2] allowing experiment to start from an intermediate stage rather than empty execution state. Resources were only examined related to availability, utilization and count. A hybrid approach for business process modeling and simulation with limited resources modeling related to concurrency aspects were addressed in [27]. The approach examined concurrent execution of tasks through resources multitasking to avoid blocking resources for long time. Our work is considered complementary to both approaches. However, the concurrent execution by human resources is of very limited use especially when it comes to knowledge intensive tasks.

The work in [28] discusses supporting business process simulation model construction by using event logs. They provide means to extract Resources behavior from logs. The modeling of resources behaviors was supported by roles, schedules, handling procedures, resources unavailability and grouping resources based on similar activities handling. Yet, constraints on resource behavior was not identified by the authors. Thus, our work complements theirs in that regard.

L-Sim, a BPMN-based simulator was presented in [6]. L-Sim supports a subset of pull and push patterns. OXProS - an Open and Extensible Process Simulator for BPMN was presented in [1], the tool supports modeling role-based and chained execution distribution in workflow resources patterns. Other resource patterns were left for future work. A blueprint architecture was proposed in [29] for a business process simulation engine that focuses on BPMN models. Resources modeling was only concerned with resource availability. Authors in [30] presents an open source and extensible BPMN process simulator that provide basic resource support provided by BPMN. Resources are defined globally using a configuration file provided in XML to be involved in multiple processes.

A survey on business process simulation tools was conducted in [4]. Amongst the evaluation criteria was the ability to model the resources perspective. Amongst the evaluated tools were Arena² (a general-purpose simulation tool) and CPN tools³. The study results indicated that CPN tools provides better support for resources representation than Arena. However, both require quite an effort to model resources behavior in a way that might be inaccessible to business process experts.

² <https://www.arenasimulation.com/>.

³ <http://cpntools.org/>.

5 Conclusion and Outlook

In this paper, we have discussed the importance of workflow resource patterns to reflect resources representation and utilization. Here, we refined the relationships among workflow resource patterns using *complement-perspective* to build a minimal set of constructs that can be reused in the related patterns. We reviewed the limitations of modeling resource perspective within BPSim standard V2.0 and we took a first step towards a resource aware BPSim extension (RBPSim) combining BPSim standard with workflow resource patterns. RBPSim provides a tool independent exchange format for so-called simulation scenarios including resources perspective.

In future work, we plan to introduce the remaining resource patterns. Moreover, we aim to start implementing the extended BPSim metamodel. Implementation may have two directions, the first is to seek an open source BPS tool that supports BPSim standard to apply the extended metamodel. The second direction is to implement the extension using a general-purpose simulation tool (e.g. ExtendSim).

References

1. Garcia-Banuelos, L., Dumas, M.: Towards an open and extensible business process simulation engine. In: CPN Workshop (2009)
2. Wynn, M.T., Dumas, M., Fidge, C.J., ter Hofstede, A.H.M., van der Aalst, W.M.P.: Business process simulation for operational decision support. In: ter Hofstede, A., Benatallah, B., Paik, H.-Y. (eds.) BPM 2007. LNCS, vol. 4928, pp. 66–77. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-78238-4_8
3. van der Aalst, W.M.P.: Business process simulation revisited. In: Barjis, J. (ed.) EOMAS 2010. LNBIP, vol. 63, pp. 1–14. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-15723-3_1
4. Jansen-Vullers, M., Netjes, M.: Business process simulation-a tool survey. In: Workshop and Tutorial on Practical Use of Coloured Petri Nets and the CPN Tools, vol. 38, pp. 1–20 (2006)
5. Stroppi, L.J.R., Chiotti, O., Villarreal, P.D.: A BPMN 2.0 extension to define the resource perspective of business process models. In: XIV Congreso Iberoamericano en Software Engineering (2011)
6. Waller, A., Clark, M., Enstone, L.: L-SIM: Simulating BPMN diagrams with a purpose built engine. In: Simulation Conference, pp. 591–597. IEEE (2006)
7. Stroppi, L.J.R., Chiotti, O., Villarreal, P.D.: Defining the resource perspective in the development of processes-aware information systems. *Inf. Softw. Technol.* **59**, 86–108 (2015)
8. Nakatumba, J., Rozinat, A., Russell, N.: Business process simulation: how to get it right. In: International Handbook on Business Process Management. Springer (2009)
9. Russell, N., van der Aalst, W.M.P., ter Hofstede, A.H.M., Edmond, D.: Workflow resource patterns: identification, representation and tool support. In: Pastor, O., Falcão e Cunha, J. (eds.) CAiSE 2005. LNCS, vol. 3520, pp. 216–232. Springer, Heidelberg (2005). https://doi.org/10.1007/11431855_16

10. Freitas, A.P., Pereira, J.L.M.: Process simulation support in BPM tools: The case of BPMN. In: Proceedings of 2100 Projects Association Joint Conferences (2015)
11. WfMC: BPSim - business process simulation specification (2016)
12. OMG: Business process model and notation (BPMN version 2.0) (2011)
13. Laue, R., Müller, C.: The business process simulation standard (BPSim): Chances and limits. In: ECMS, pp. 413–418 (2016)
14. Affi, N., Awad, A., Abdelsalam, H.M.: Rbpsim: A resource-aware extension of BPSim using workflow resource patterns. In: 10th Zeus Workshop, vol. 2072, Dresden, Germany, pp. 31–38. CEUR-WS (2018)
15. Russell, N., Ter Hofstede, A.H., Edmond, D., van der Aalst, W.M.: Workflow resource patterns. Technical report, BETA Working Paper Series, WP 127, Eindhoven University of Technology, Eindhoven (2004)
16. Russell, N., van der Aalst, W.M., ter Hofstede, A.H.M.: Workflow Patterns: The Definitive Guide. The MIT Press (2016)
17. Mendling, J., Ploesser, K., Strembeck, M.: Specifying separation of duty constraints in BPEL4People processes. In: Abramowicz, W., Fensel, D. (eds.) BIS 2008. LNBIP, vol. 7, pp. 273–284. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-79396-0_24
18. Tumay, K.: Business process simulation. In: Winter Simulation, pp. 93–98. IEEE Computer Society (1996)
19. Hlupic, V., Robinson, S.: Business process modelling and analysis using discrete-event simulation. In: Winter Simulation Conference, pp. 1363–1370. IEEE Computer Society Press (1998)
20. Mathew, B., Mansharamani, R.: Simulating business processes—a review of tools and techniques. *Int. J. Model. Optim.* **2**(4), 417 (2012)
21. Cabanillas, C., Resinas, M., Ruiz-Cortés, A.: RAL: a high-level user-oriented resource assignment language for business processes. In: Daniel, F., Barkaoui, K., Dustdar, S. (eds.) BPM 2011. LNBIP, vol. 99, pp. 50–61. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-28108-2_5
22. Stroppi, L.J.R., Chiotti, O., Villarreal, P.D.: Extending BPMN 2.0: method and tool support. In: Dijkman, R., Hofstetter, J., Koehler, J. (eds.) BPMN 2011. LNBIP, vol. 95, pp. 59–73. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-25160-3_5
23. Awad, A., Grosskopf, A., Meyer, A., Weske, M.: Enabling Resource Assignment Constraints in BPMN. Hasso Plattner Institute, Potsdam (2009)
24. Januszczak, J., Hook, G.: Simulation standard for business process management. In: Proceedings of the 2011 Winter Simulation Conference (WSC), pp. 741–751. IEEE (2011)
25. van der Aalst, W.M.P., Nakatumba, J., Rozinat, A., Russell, N.: Business process simulation. In: Brocke, J., Rosemann, M. (eds.) Handbook on Business Process Management, vol. 1. International Handbooks on Information Systems. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-00416-2_15
26. van der Aalst, W.M.P.: Business process simulation survival guide. In: vom Brocke, J., Rosemann, M. (eds.) Handbook on Business Process Management, vol. 1. International Handbooks on Information Systems. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-642-45100-3_15/
27. Vasilecas, O., Smaizys, A., Rima, A.: Business process modelling and simulation: hybrid method for concurrency aspect modelling. *Baltic J. Mod. Comput.* **1**(3–4), 228–243 (2013)
28. Martin, N., Depaire, B., Caris, A.: The use of process mining in business process simulation model construction. *Bus. Inf. Syst. Eng.* **58**(1), 73–87 (2016)

29. Krumnow, S., Weidlich, M., Molle, R.: Architecture blueprint for a business process simulation engine. In: EMISA, vol. 172, pp. 9–23 (2010)
30. Pufahl, L., Weske, M.: Extensible BPMN process simulator. In: BPM Demo, vol. 1920. CEUR Workshop Proceedings. CEUR-WS.org (2017)



Towards Implementing REST-Enabled Business Process Choreographies

Adriatik Nikaj^(✉), Marcin Hewelt, and Mathias Weske

Hasso Plattner Institute, University of Potsdam, Potsdam, Germany
{adriatik.nikaj,marcin.hewelt,mathias.weske}@hpi.de

Abstract. When it comes to the interaction on the Web, one of the most adopted architectural styles is REST. On the modeling side, business process choreographies model the inter-organizational processes that are performed by business participants to reach a common goal. Bridging the gap between the modeling level and the RESTful interactions is challenging. We tackle this challenge by introducing an intermediary RESTful service that assures the correct execution of the choreography by guiding the participants to adhere to the commonly agreed interactions. Moreover, this service allows human participants to partake in the choreography without having a complex system in place.

Keywords: Business process choreographies · RESTful interactions
Process orchestration

1 Introduction

In the era of Internet, enterprises can focus on their core capabilities and out-source supporting activities to other business partners forming so called networked enterprises. This is achieved through fast and reliable inter-organizational communication. When it comes to the interaction on the Web, one of the most adopted architectural styles is Representational State Transfer (REST) [1]. On the modeling side, business process choreographies model the inter-organizational processes that are performed by business participants to reach a common goal. In particular, choreography diagrams are introduced as part of BPMN 2.0 [2] to formally specify the ordering of message exchanges between two or more participants.

However, there is a large conceptual gap between the modeling level and the concrete RESTful interaction. The gap can be bridged by translating the global perspective, from which the interactions between participants are modeled, to the RESTful services that participants need to provide for implementing the specified interactions. In [3], the authors partially address this gap by annotating choreography diagrams with REST implementation-specific information in a semi-automatic fashion. This hybrid model is called RESTful choreography. In this paper, we further close the gap between RESTful choreography and its

REST implementation by introducing RESTful Choreography Guide (ChoreoGuide) - a central RESTful service acting as a intermediary that enforces the correct execution of the choreography between the business actors.

ChoreoGuide is derived systematically from the RESTful choreography and the Choreography Resource Model - a commonly agreed data model with constraints on the exchange of REST resources. ChoreoGuide per se is a business process orchestration exposed via a REST interface and can be deployed on a RESTful business process engine to be executable. That means that participants do not necessarily require a process engine or a complex system in place to interact with it. We benefit from the separation of concerns, i.e. all participants are clients (from the REST perspective) to the same server (ChoreoGuide). This way the choreography platform can evolve independently from the clients underlying systems as long as the REST interface is kept unchanged. This is the case when the platform can be used by third parties as a starting point to create new business models that bridge the interaction between business actors and clients like Airbnb¹ and Easychair².

2 Preliminaries

This section introduces choreography diagrams and RESTful choreography along with a running example used for motivating our approach.

Business process choreography [2] is a specification language for modeling interactions between two or more participants. It abstracts from the internal logic of participants' own process models and provides a global perspective without a preferential viewpoint. Figure 1 describes the interaction of three participants involved in a purchase process. The skateboard manufacturer sends an order to the ball bearings supplier, which the supplier can confirm or reject. In case of the former, the supplier sends the ball bearings to the skateboard manufacturer as requested in the order. Once the ball bearings are delivered the skateboard manufacturer determines the percentage of defective units. If the percentage is lower than 5% the manufacturer initiates a payment request to the payment organization (the third participant). Otherwise, the skateboard manufacturer sends back the ball bearings leading to an unsuccessful purchase.

The main composing unit of a choreography diagram is the choreography task (the rounded rectangles in Fig. 1), which represent either sending a single message or a send message followed by a reply message. In a single choreography task, the participant who sends the message is called the initiator and the message is referred to as initiating message, while the recipient may send back an optional response message. Similarly to activities in business processes, the choreography tasks are ordered via sequence flows. To express parallel or exclusive branches a subset of business process gateways is used. In this paper we consider only three type of gateways: parallel, exclusive and event based gateways.

¹ <https://www.airbnb.com/>.

² <http://www.easychair.org/>.

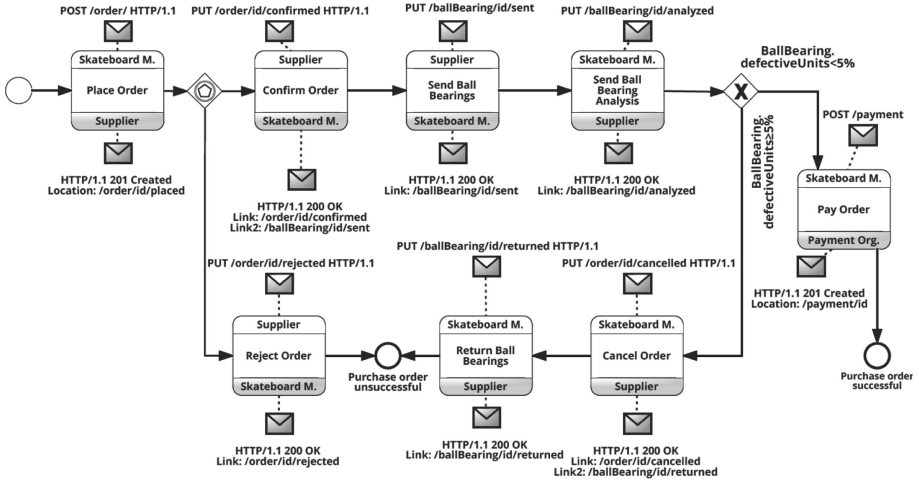


Fig. 1. RESTful choreography for the purchase of ball bearings

REST [1] is an architectural style for web services. Services provide their resources to clients via unique identifiers (URI) allowing only a specified set of operations (REST verbs) with a stateless nature. Key feature of RESTful interactions is the usage of hypermedia as the engine of application state (HATEOAS). Links provided in the response guide the clients through subsequent interactions. These architectural constraints contribute to a simplified system architecture, better performance and improved evolvability of services without affecting the interaction.

In order to bridge the conceptual gap between process choreographies and their implementation as RESTful interactions, [4] introduces RESTful choreography diagrams. They are a light-weight extension of BPMN choreography diagrams with REST annotations (see Fig. 1). Choreography tasks are enhanced to REST tasks, in which the initiating message is annotated with a REST request method (e.g. *POST /order*) and the return message is annotated with response status codes (e.g. *201 Created*, *Location: /order/id*). Hyperlinks can be added in the request and response to guide the participants, who can follow them according to the behavior modeled in the RESTful choreography. Moreover, in [3] it is shown that the generation of REST annotations can be automatized to a good extend.

3 RESTful Choreography Guide

In this section we describe RESTful Choreography Guide as a blueprint for developing RESTful platforms that implement choreographies. ChoreoGuide is a process model orchestration with REST implementation-specific information that can be deployed in any RESTful business process engine. We purposely

choose a technology-agnostic approach that can be adopted and applied to different technologies. However, for evaluating the feasibility of our approach, we provide an implementation architecture for deploying ChoreoGuide onto a concrete platform (see Sect. 4).

Since the two main artifacts treated by the choreography diagram are the messages and their order, RESTful Choreography Guide has two functional requirements:

1. Validate the message payload
2. Orchestrate the control flow of the message exchanges

In a RESTful choreography, the annotation of an initiating message contains a REST verb, the name of the resource and its state, e.g., *PUT order/id/-confirmed*. In order to check the validity of the resource and its state we need more information about the data being exchanged than provided by the choreography. For that, the participants need to agree on a common choreography resource model (see Sect. 3.1). This model is required by ChoreoGuide to check the validity of the resource and its state. The choreography resource model is specified using two standards: UML class diagram and object constraint language (OCL) [5].




Fig. 2. Approach overview

Regarding the second requirement, the control flow and the message exchanges in a RESTful choreography are modeled from a global perspective. We provide a mapping, detailed in Sect. 3.2, of the choreography control flow to business process control flow where message exchanges are implemented by process orchestration activities.

Figure 2 provides an overview of our approach for implementing the RESTful choreography via ChoreoGuide. Our approach takes as input a RESTful choreography and its resource model. The target model is a BPMN [2] process model with embedded REST information that can be deployed to a business process engine capable of providing a RESTful API.

3.1 Choreography Resource Model

In this section, we describe the choreography resource model. The resource model is represented in part as an UML class diagram (see Fig. 3) that serves the

purpose of specifying the main REST resources, their state and their relation to each other and to choreography participants as well. It is highly influenced by the domain area and specified by the involved participants alongside the RESTful choreography diagram at design time.

Figure 3 illustrates the static view of the choreography resource model for our running example modeled in Eclipse Modeling Framework³. Every *ResourceModel*, independently of the use case, is always composed of at least two *Participants* and at least one *RestResource*. Depending on the use case, participants and REST resources are added as classes that inherit the abstract classes *Participant* and *RestResource* respectively. The attributes and the associations are filled depending on the use case domain. For example, *Order* has a *creationDate* and *BallBearing* as an *item*.

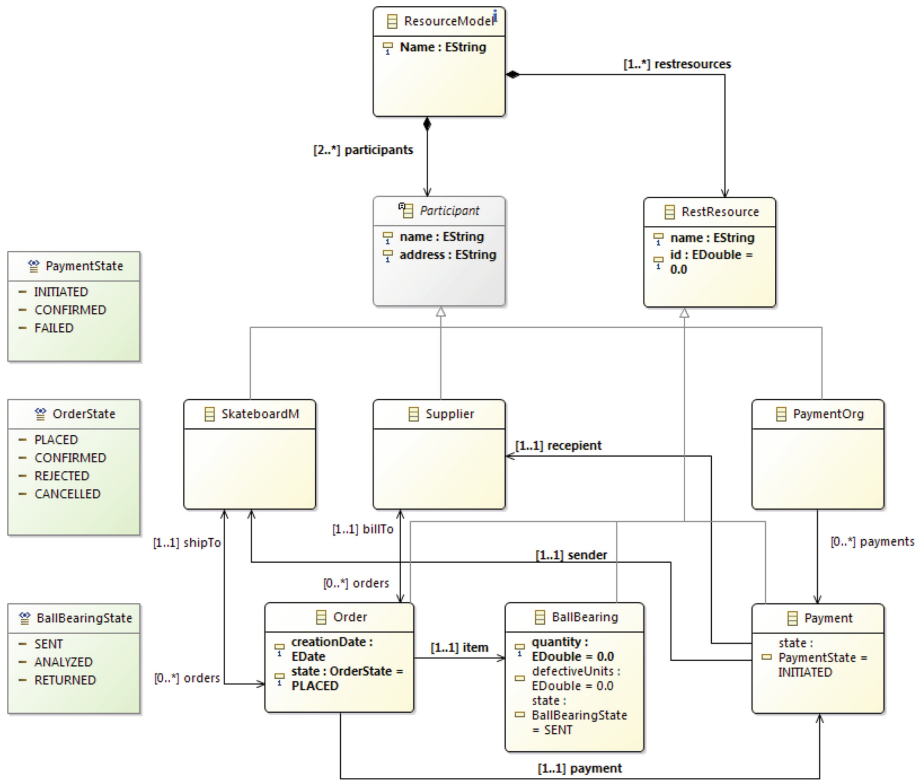


Fig. 3. Choreography Resource Model (static model)

However, a static model of the resources is not sufficient. Resource states are not mere attributes that can be assigned by participants. States represent

³ <https://eclipse.org/modeling/emf/>.

certain conditions of the resource which are aggregated from the concrete value assignments of the resource attributes. For example, resource *order* is in the state *placed* only if the quantity of the ball bearings is greater than 0. We use OCL [5] to formally express such conditions that must hold during the entire execution of the choreography. Below is the respective OCL expression for the aforementioned condition.

```
context order
inv nonEmptyOrder: self.state = OrderState : : PLACED
implies self.item.quantity > 0
```

The OCL expression refers to the UML class diagram. In this expression we look at the class *Order*. Then we identify an invariant (*nonEmptyOrder*) that must hold during the instantiation of the *order* objects. The invariant in this case is an implication where the premise is the resource in state *PLACED* and the consequence is an expression that must hold for that state. In fact, this is the pattern we propose to check for any resource object the validity of its state:

```
context <<resource>>
inv <<resource state>>
implies <<condition>>
```

In order to use the OCL expression, each *RESTResource* child class must have an attribute *state* of the type *Enumeration* that contains all the possible states of that particular resource, e.g., *Payment* can be in state initiated, confirmed, or failed. The condition of resource state can be arbitrary complex and long (to the extend allowed by OCL). OCL provides the possibility to navigate the class diagrams and express complex relations between resources. Using a standard like OCL allows the developers to choose the desired tools that support evaluating the expressions. Another key benefit is that OCL is free from side effects. This means that checking whether an OCL expression holds does not affect the running program. When it comes to choreographies, this is very useful because ChoreoGuide should only reply whether the state change request is valid or not without interfering with the internal logic of the participants' process.

3.2 From RESTful Choreography to Process Orchestration

In this section we present how the process orchestration is systematically derived from the RESTful choreography and the resource model. We consider the translation of the RESTful choreography's tasks and three types of gateways: parallel gateway; data-based exclusive gateway; and event-based gateway as the most commonly used choreography elements. In the following, these choreography elements are mapped to business process constructs. After that, all the individual derived constructs are concatenated to form the whole target business process.

In a RESTful choreography, there are four types of REST Tasks: GET; DELETE; POST; and PUT. The latter two represent a resource creation or a resource state change. These two kinds of tasks are translated to the following

business process constructs (see Fig. 4): a receive task; a service task; an exclusive gateway which branches the control flow in two parts; a “fail” message intermediate throwing event; a “successful” message intermediate throwing event; and a send task that forwards the initiator’s request. Specifically, as illustrated in Fig. 4, the receiving task is accessed for the participants by the REST request specified in the REST task, e.g., *POST /resource/ HTTP/1.1*. In case of a POST or PUT a resource or a representation of it is delivered as a message payload to the server. With POST, a new resource is created in an initial state and PUT is used to update a resource with a new state. After the resource is received from the server, an automatic check is performed in the subsequent step by a service task. This task automatically checks whether or not the state-change request is valid by evaluating the respective OCL expression in the resource model. In case the OCL expression is evaluated to false, the server replies to the client (the participant who sent the request) with a *HTTP/1.1 428 Precondition Required*. Otherwise, the server replies with a *HTTP/1.1 201 Created* (for POST) or *HTTP/1.1 200 OK* (for PUT). RESTful choreography diagrams model only the successful interaction between the client and the server. Hence, the response of a valid request used in the process orchestration is derived from the choreography. Finally, the link for getting the new resource state is forwarded to the recipient.

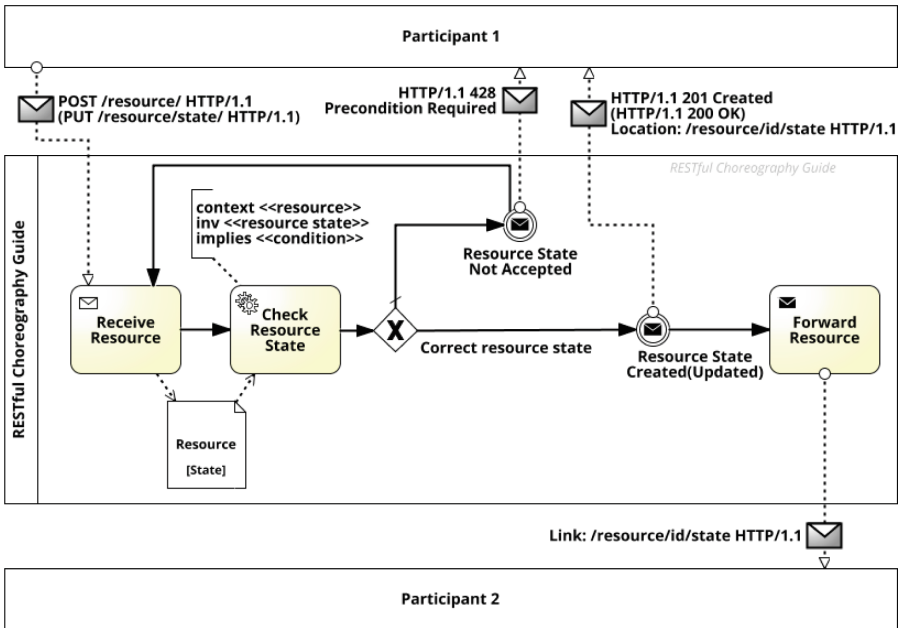


Fig. 4. Business process construct for POST and PUT (in brackets) REST tasks

GET is used to read the state of the resource and DELETE to remove the resource from the location specified in the URI. In case of GET task, the state of the resource does not change and, therefore, there are no conditions to be checked. The corresponding process orchestration (see Fig. 5) has three consecutive nodes linked by sequence flows: a receiving task for the request, a message intermediate throwing event for replying to the initiator by sending the resource (or a partial representation of it) and a send message task for notifying the recipient that the resource has been read.

When a DELETE request is sent, there are no conditions to be checked (similarly to GET task) in terms of attribute values. Hence, the corresponding process orchestration is the same as in the GET case (see Fig. 5), but a “resource deleted” message notification is sent back in lieu of a reply containing the resource. It is worth mentioning that the delete request is not forwarded to the participant prior to the server reply. This is because this delete request is part of the intended behavior specified by the RESTful choreography. If any participant would arbitrary request to delete a resource, ChoreoGuide will not accept the request unless it is part of the designed behavior.

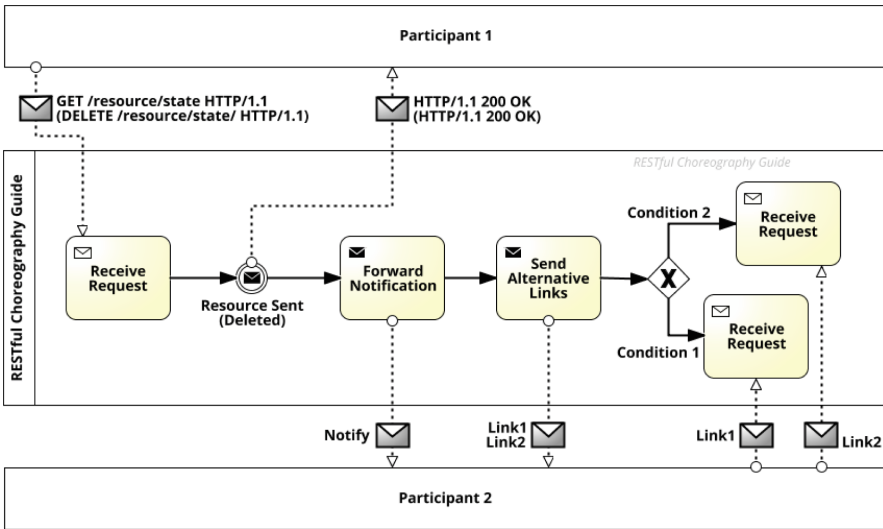


Fig. 5. Business process construct for GET and DELETE (in brackets) REST tasks followed by the construct for exclusive gateway

Choreography’s parallel gateway with m outgoing sequence flows is mapped to the following business process construct: a send task that sends a message containing n links where n is the number of choreography tasks that immediately follow the parallel gateway. The send task is followed by a process parallel gateway (same syntax with the choreography’s parallel gateway) with the m

outgoing sequence flows. n equals m when the choreography parallel gateway is immediately followed by only REST tasks.

Similarly to the parallel gateway, the choreography's event-based gateway with m outgoing sequence flows is mapped to a send task followed by a process event-based gateway with m outgoing sequence flows. The send task sends a message containing n links, where n is the number of REST tasks that immediately follow the gateway. A concrete example is depicted in Fig. 7.

The choreography's exclusive gateway follows the exact mapping as event-based gateway - the target process construct consists of a send task followed by process exclusive gateway (see Fig. 5). However in this case, the outgoing sequence flows are conditional sequence flows which also need to be mapped accordingly. This means that the conditions need also to be evaluated from ChoreoGuide. According to the specification of choreography's exclusive gateway, the data upon which the conditions are evaluated is passed previously to the participants affected by the gateway (participants who immediately follow the gateway). For example, the number of defective units is passed from the skateboard manufacturer to the supplier via the ball bearings analysis. This value is found in the instance of the resource model (see Fig. 3). Since the resource model is managed by ChoreoGuide, it has sufficient data to evaluate the conditional flows. This constitutes an additional solution that solves the classical problem of implementing the choreography exclusive gateway [6].

Finally, the choreography start events, end events and join gateways are mapped accordingly to ChoreoGuide's business process. The concatenation of all generated business process constructs is realized by connecting them with a single sequence flow as defined in the RESTful choreography. After the concatenation, a single reduction rule is applied: For every gateway in the RESTful choreography that immediately follows a REST task, the Forward Resource/Notification task is merged with the send task preceding the gateway (the task sending the links) if the message recipient is the same (like in Fig. 5). The payloads of the outgoing messages (of the send tasks before the merge) are added up and form the new message's payload. A snippet of the resulting mapping from our running example is depicted in Fig. 7. The Forward Order send task is merged with the send task that immediately precedes the event-based gateway. The resulting task, besides notifying the supplier about the order placement, sends the links needed to execute the upcoming event-based gateway.

4 Implementation Architecture

This section presents the architecture of the proposed approach and discusses how the generated ChoreoGuide orchestration process is deployed to an existing process engine to validate the feasibility of the approach. Figure 6 shows the main components of the architecture, (1) the parser responsible for parsing the RESTful choreography and resource model, (2) the generator that derives the ChoreoGuide orchestration process and deploys it, and (3) the process engine

Chimera⁴. Chimera is an academic process engine, developed to validate research approaches in the field of BPMN and case management. It is a good fit for RESTful choreographies, because it exposes running processes via a RESTful API. As free and open source software it is easily extensible for the purpose of this contribution.




Fig. 6. Architecture of ChoreoGuide

The Chimera engine requires process model to define the data classes it operates on by providing a data model. For ChoreoGuide we use the sub-classes of **Resource** from the resource model of the choreography to be deployed. The choreography resource model is used to define data objects and their attributes.

The Addressing Problem. As part of the mapping to a concrete execution engine the addressing of ChoreoGuide and participants has to be addressed. Since participants will conduct a choreography multiple times, the engine needs to ensure that messages are correlated to the correct *instances*. Additionally, the Chimera RESTful API makes data objects available only in the context of a process instance. Therefore the URI annotations used in the RESTful choreography, e.g. */resource*, need to be prefixed by the following parts: (1) base URI where the engine is reachable in the network, (2) identifier of the choreography, (3) identifier of the choreography instance. The resulting URI for an order resource thus might look like <https://example.org/chimera/api/choreography/5/instance/23/order>. In case that one choreography instance handles multiple resources of the same type, e.g. multiple orders, an identifier of the concrete order needs to be appended to the URI.

One of the benefits of our approach is that participants are free to implement their part of the choreography in different ways. Hence, the ChoreoGuide needs different methods to notify them. A simple email suffices as vessel for hyperlinks. In our example the supplier receives an email with a short order description and two hyperlinks (see Fig. 7) to either reject or confirm the order. Depending on the systems used by the participants other push technologies could be used as well [7]. However, we leave the details of such integration for future work.

When the initiator of a the first choreography task sends the initial message no instance of the ChoreoGuide orchestration exists yet. Therefore, the first message needs to be translated to a POST to the Chimera engine for instantiating the ChoreoGuide instance. The last point we want to discuss is how the

⁴ <https://bptlab.github.io/Chimera>.

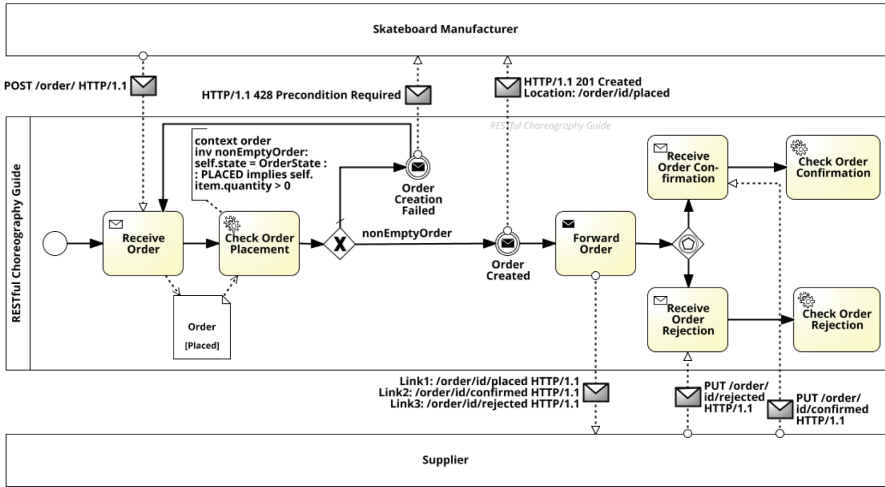


Fig. 7. RESTful Choreography Guide sample of the running example

OCL conditions are checked in the process engine. The OLC conditions refer to attributes of resources defined in the resource model. This corresponds to the data objects managed by the ChoreoGuide instance. Therefore, the conditions are used as annotations on the sequence flows following the exclusive gateway. When they evaluate to true the choreography continues, otherwise a response with HTTP status code 428 is send to the requester.

5 Related Work

There is a plethora of related work covering the execution of the choreography's control flow. Zaha et al. are among the first authors to put forward the issue of local enforceability in [8] in the context of web services standards like SOAP and WSDL [9]. They propose an approach for analyzing if the relation between service interactions described from a global perspective could be enforced locally by generating BPEL [10] templates.

In [11], Barros et al. focus on the correlation between the exchanging messages. The correlation between messages is an important factor when it comes to the implementation of choreographies, and thus, to their enforceability. They introduce a framework for classifying correlation scenarios in the context of SOA and argue that BPEL does not support certain patterns. Decker et al. [12] extend BPEL web service composition standard [10] to bring process orchestration closer to choreographies. BPEL4Chor implements process choreographies by composing existing BPEL service orchestrations. This bottom-up approach is based on web services standards like SOAP and WSDL.

In these related work the local enforceability is checked from the control flow perspective. The authors assume that the data used in the choreography is

interpreted equally. In our approach, having a common resource model designed by the participants improves the interpretability of the inter-organizational data. Moreover, the related work is about relating choreographies to SOAP and WSDL. Our work, instead, focuses on relating business process choreographies to RESTful services, which have seen prominent uptake in the last years.

An approach that focus on the relation between BPMN and REST is provided in [13]. The authors provide an activity-centric mapping of business process elements to REST interfaces. In contrast, ChoreoGuide’s business process provides a data-centric mapping to REST interfaces driven by the messages exchange defined in the choreography. In our approach, we unify REST resources with business process data objects and enable the transitions of their states throughout the choreography execution.

At last, a related work that employs an intermediary platform for executing the choreographies is introduced by Weber et al. [14]. The proposed platform, however, is based on blockchain - a global infrastructure capable of executing programs called *smart contracts*. The goal is to ensure a correct execution of choreographies in the presence of untrusted participants. Our platform assumes a higher trust level between the participants since they have to agree on the choreography and the resource mode. Once ChoreoGuide is deployed, the participants can monitor the state of the choreography at real-time and interact with it only as allowed by ChoreoGuide. Additionally, executing the derived business orchestration on a blockchain is more costly than using, for example, a cloud service [15].

6 Conclusion

In this paper we introduce a novel approach towards implementing business process choreographies. Existing work, focus on deriving public processes for each participant and enforce their execution. We propose a central RESTful service that takes the participants “by the hand” and guides them through every step along the choreography. The service employs the REST architectural style to take advantage of the REST constraints.

One key benefit of our approach is that the participants are not required to run a process engine or a complex system to interact with each other. This is particularly important for the interaction with human clients that can participate in the choreography via a simple web browser. Our approach fully employs the principle of HATEOAS, i.e. hyperlinks are indeed the engine of the choreography state.

Another important benefit is that ChoreoGuide not only enforces the control flow but validates the correctness of the resource state changes using the choreography resource model. In addition, having such a model helps the participant to agree on a common data structure and avoid problems coming from the misinterpretation of the inter-organization data. Future work will focus on the automatic deployment of ChoreoGuide into a specific business process engine.

References

1. Fielding, R.T.: Architectural Styles and the Design of Network-based Software Architectures. PhD thesis, AAI9980887 (2000)
2. OMG: Business Process Model and Notation (BPMN), Version 2.0, January 2011. <http://www.omg.org/spec/BPMN/2.0/>
3. Nikaj, A., Weske, M., Mendling, J.: Semi-automatic derivation of restful choreographies from business process choreographies. *Softw. Syst. Model.*, January 2018
4. Nikaj, A., Mandal, S., Pautasso, C., Weske, M.: From choreography diagrams to RESTful interactions. In: Norta, A., Gaaloul, W., Gangadharan, G.R., Dam, H.K. (eds.) *ICSOC 2015*. LNCS, vol. 9586, pp. 3–14. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-50539-7_1
5. OMG: Object Constraint Language (OCL), Version 2.0, May 2006. <http://www.omg.org/spec/OCL/2.0/>
6. Nikaj, A., Batoulis, K., Weske, M.: REST-enabled decision making in business process choreographies. In: Sheng, Q.Z., Stroulia, E., Tata, S., Bhiri, S. (eds.) *ICSOC 2016*. LNCS, vol. 9936, pp. 547–554. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46295-0_34
7. Pautasso, C., Wilde, E.: Push-enabling RESTful business processes. In: Kappel, G., Maamar, Z., Motahari-Nezhad, H.R. (eds.) *ICSOC 2011*. LNCS, vol. 7084, pp. 32–46. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-25535-9_3
8. Zaha, J.M., Dumas, M., Hofstede, A.T., Barros, A., Decker, G.: Service interaction modeling: bridging global and local views. In: 10th IEEE International Enterprise Distributed Object Computing Conference, EDOC 2006, pp. 45–55. IEEE (2006)
9. Alonso, G., Casati, F., Kuno, H., Machiraju, V.: *Web Services*. Springer, Heidelberg (2004). <https://doi.org/10.1007/978-3-662-10876-5>
10. Jordan, D., Evdemon, J., Alves, A., Arkin, A., Askary, S., Barreto, C., Bloch, B., Curbera, F., Ford, M., Golland, Y., et al.: *Web services business process execution language version 2.0*. OASIS standard 11, 10 (2007)
11. Barros, A., Decker, G., Dumas, M., Weber, F.: Correlation patterns in service-oriented architectures. In: Dwyer, M.B., Lopes, A. (eds.) *FASE 2007*. LNCS, vol. 4422, pp. 245–259. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-71289-3_20
12. Decker, G., Kopp, O., Leymann, F., Weske, M.: Bpel4chor: Extending BPEL for modeling choreographies. In: *IEEE International Conference on Web Services, ICWS 2007*, pp. 296–303. IEEE (2007)
13. Pautasso, C.: BPMN for REST. In: Dijkman, R., Hofstetter, J., Koehler, J. (eds.) *BPMN 2011*. LNBIP, vol. 95, pp. 74–87. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-25160-3_6
14. Weber, I., Xu, X., Riveret, R., Governatori, G., Ponomarev, A., Mendling, J.: Untrusted business process monitoring and execution using blockchain. In: La Rosa, M., Loos, P., Pastor, O. (eds.) *BPM 2016*. LNCS, vol. 9850, pp. 329–347. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-45348-4_19
15. Rimba, P., Tran, A.B., Weber, I., Staples, M., Ponomarev, A., Xu, X.: Comparing blockchain and cloud services for business process execution. In: *2017 IEEE International Conference on Software Architecture, ICSA 2017*, Gothenburg, Sweden, 3–7 April 2017, pp. 257–260 (2017)



Disambiguation of DMN Decision Tables

Kimon Batoulis^(✉) and Mathias Weske

Hasso Plattner Institute, University of Potsdam, Potsdam, Germany
{Kimon.Batoulis,Mathias.Weske}@hpi.de

Abstract. The Decision Model and Notation (DMN) is a specification for the design of decision models. The logic of those decisions can be expressed in standardized decision tables. The standard allows to create tables that contain overlapping rules. This leads to ambiguities and we argue that such tables are hard to understand and unsuitable for analysis tasks. This paper describes an algorithm to transform ambiguous DMN decision tables to equivalent ones that only contain exclusive rules, thereby resolving ambiguities. We implemented our algorithm and evaluated it against a set of synthetic decision tables.

Keywords: DMN · Decision table analysis · Disambiguation

1 Introduction

In the Decision Model and Notation (DMN) specification, decision tables are a standardized way of expressing decision logic within a decision model [8]. Therefore, the majority of decision models consist of tables and much research has been conducted based on this representation, such as analysis and simplification [5, 6], mining [3], and verification of their integration with process models [1, 2].

In principal, decision tables provide a clear view of how decisions are taken because of its tabular representation of rules that map inputs to outputs. In the simplest case, for a given input, at most one rule matches and its output is returned. However, the DMN standard allows to design more complex tables where for a given input more than one rule matches. In these cases, there are standardized possibilities of resolving such conflicts. For example, one could just return the outputs of all matching rules. The problem with such tables is that their input-output behavior is not clearly visible anymore. This is because each input could match any number of rules, whose outputs are then aggregated in some way.

In this paper, we describe an algorithm that transforms any type of DMN decision table that uses S-FEEL syntax into a behaviorally equivalent table with exclusive rules, i.e., a table that has the same input-output behavior but for which a given input matches at most one rule. Such tables are *unambiguous* because they have a clear input-output behavior and are therefore more comprehensible and also better suited for analysis tasks such as the ones described

in [2,5]. We implemented our disambiguation procedure in the Camunda *dmn-js* editor and evaluated its performance against a synthetic decision table data set.

The remainder is structured as follows. We motivate our approach by introducing a running example and laying the foundations for our work in Sect. 2. Section 3 describes the table disambiguation procedure, which is evaluated empirically in Sect. 4. Section 5 is devoted to related work, and Sect. 6 concludes the paper.

2 Motivation and Prerequisites

This section first gives an example of a decision table that motivates the benefits of decision table disambiguation. Afterwards, necessary details of DMN decision tables are described and existing techniques for their analysis are explained.

2.1 Motivational Example

Figure 1 shows a DMN decision table. It has two inputs, *Income* and *Assets* (given in thousands), and one output *Credit Rating*. Both, inputs and outputs are associated with a type and an optional restriction on the allowable values of that type. For instance, the inputs in the example are real numbers which are greater than or equal to zero and the output is a string, but it can only take on the values that are listed below its name.

R	Income (k) <i>Number ≥ 0</i>	Assets (k) <i>Number ≥ 0</i>	Credit Rating <i>A, B, C, D, E</i>
1	≤ 30	≤ 30	A
2	∈ [10..60]	∈ [10..25]	B
3	∈ [20..95]	∈ [40..90]	C
4	≥ 80	-	D
5	≥ 40	≥ 85	E

Fig. 1. A DMN decision table with overlapping rules and a *rule order* hit policy

The table has five rules, each of which matches for certain combinations of input values and relates them to an output value. For example, rule 1 matches for the input (15, 10) and relates it to the output *A*. However, the same input is also matched by rule 2, which relates it to a different value, namely *B*. This means that rules 1 and 2 are overlapping, leading to ambiguities regarding the input values for which they both match: Which output value should be chosen for these input values?

Such ambiguities are resolved by the hit policy of the decision table, represented by a single letter in its upper left corner. For example, the letter *R* in the upper left corner in Fig. 1 means that this table has a *rule order* hit policy. Under this policy, ambiguities due to multiple matching rules are resolved by just putting the outputs of all the matching rules in a list sorted by the order

of those rules in the table. Therefore, the table would actually relate the input (15, 10) to the output $[A, B]$.

This leads to the problem of determining the input-output behavior of a DMN decision table. More precisely, what are the possible output values of the table, and which input values generate which of these output values? For example, the possible output values for the table in Fig. 1 are the following:

$$A, B, C, D, E, [A, B], [C, D], [C, E], [D, E], [C, D, E],$$

which is a far from obvious by just looking at the table. Similarly, it is not obvious which of these output values is generated from given input values. For instance, given the input (65, 90), one can see that it matches rule 3, but since this is a table with overlapping rules, one cannot be sure that this is the only matching rule. One has to keep scanning the rules until all matches are found—an error-prone process. In fact, also rule 5 matches for the input and the output is $[C, E]$.

To summarize, DMN decision tables with overlapping rules are unintuitive and hard to understand. Moreover, because of their obscure input-output behavior they are unsuited for analysis tasks such as checking their correct integration with business process models [1, 2, 7]. Hence, in this paper, we propose an algorithm that transforms any type of DMN decision table into one that has only exclusive rules—a process we call disambiguation. Since the rules are exclusive, the possible output values of the table are simply given by the outputs of the individual rules—no complicated combination of multiple matching rules is required—leading to a clear input-output behavior of that table.

2.2 DMN Decision Table Hit Policies

As already mentioned above, DMN decision tables are allowed to have overlapping rules. This leads to conflicts when a given input matches more than one rule. Such conflicts are settled based on a variety of hit policies. In principle, DMN divides these policies into *single*- and *multi*-hit policies. For a given input, the former always returns the output of a single rule—even if multiple rules match—while the latter returns the outputs of all matching rules, potentially aggregated in some way.

There are four single-hit policies, only one of which guarantees that only one rule matches for a given input:

- *unique*: all rules of the table are exclusive, such that at most one rule matches for a given input.
- *any*: multiple rules can match for a given input, but they all have the same output such that there are no ambiguities.
- *first*: if multiple rules match for a given input, the rule that comes first in the decision table is chosen to return the output.
- *priority*: similar to first-hit, but in this case the rule with the highest priority is chosen, where the priority is given by the ordered list of allowable values displayed below the output's name.

The multi-hit policies are specified as follows:

- *rule order*: all matching rules' outputs returned in a list ordered by the appearance of the rules in the table.
- *output order*: all matching rules' outputs returned in a list ordered by the priority of the rules.
- *collect*: the outputs of all matching rules are collected and then aggregated in a predefined way, for example by summing up the output values or taking the maximum.

Our approach supports all of these hit policies, and translates tables with a non-unique policy to an equivalent table with a unique policy.

2.3 Decision Table Analysis

The algorithm for disambiguating decision tables presented in this paper builds on an analysis approach for finding overlapping and missing rules in a DMN decision table described in [5,6]. This approach is based on a geometric interpretation of DMN tables. Under this interpretation, the rules of a table are represented as hyperrectangles with n dimensions, where n corresponds to the number of inputs of the table. For example, the geometric interpretation of the table in Fig. 1 is given in Fig. 2 as a set of 2-dimensional hyperrectangles (i.e., rectangles).

Each rectangle corresponds to a rule of the table and is labeled with a corresponding identifier. For instance, the rectangle $r2$ represents the second rule of




Fig. 2. Decision table as a set of hyperrectangles corresponding to the rules in the table in Fig. 1

the table. This rule states that the *Income* must be between 10 and 60 and the *Assets* between 10 and 25, which is exactly the area covered by $r2$.

The problem of finding overlapping rules is now formulated as that of finding intersecting rectangles. Similarly, finding missing rules is equivalent to finding all the “white” areas that are not covered by any of the rectangles. These problems can be solved by the sweep line approach [4], which analyzes one dimension after another by sweeping a line through it. For example, Fig. 2 shows a line being swept through the *Income* dimension, from left to right. The problem of finding overlapping rules, for instance, is then solved as follows: Sweep a line through the first dimension, collecting as many rules as possible that overlap in this dimension. Then, check which of the collected rules also overlap in the other dimensions, by sweeping a line through those dimensions.

3 Table Disambiguation

This section is dedicated to our algorithm for disambiguating any DMN decision table. For any given input to a DMN decision table with a set of rules R , any subset of R can match, if the table is not a *unique-hit* table. Given the geometric interpretation of the decision table in the running example, it is quite easy to see which rule(s) match(es) for which input.

For example, for all inputs where $Income \in [10, 30]$ and $Assets \in [10, 25]$ rules 1 and 2 match. For $Income \in (30, 60]$ and $Assets \in [10, 25]$, however, only rule 2 matches. Of course, there can also be inputs for which no rules match. These are all the white “gaps” between the rectangles in Fig. 2. This means that the table is incomplete.

Based on the line sweeping technique, we developed an algorithm that will—for any DMN table—find out for which inputs which rules match. In case the table is not a *unique-hit* table, it can afterwards be translated to one that is unique. This means that the rules are then exclusive, leading to a more understandable, disambiguated decision table.

Algorithm 1 shows the pseudocode of our algorithm. The initial call to the algorithm will look like this: $findMatchingRules(R, [], 0, N, [])$, so that R is the set of rules of the table, $matchingIntervals$ is an initially empty list, i is an index initially set to 0, N is the number of inputs of the table, and $matchingRules$ is an initially empty list.

Then, the algorithm starts sweeping through the first dimension. This requires sorting the endpoints of the intervals covered by each rule in this dimension in ascending order (line 2). Then the endpoints are iterated over (lines 4–12). Whenever the current *endpoint* is a lower one (line 8), the rule belonging to this endpoint is considered active, and is therefore added to the list of *activeRules* (line 9). Contrarily, if the current *endpoint* is an upper one (line 10), the rule belonging to this endpoint is considered dead, and is therefore deleted from the list of *activeRules* (line 11).

For example, if $i = 0$, then $sortedEndpoints = [r1(0), r2(10), r3(20), r1(30), r5(40), r2(60), r4(80), r3(95), r4(\infty), r5(\infty)]$. Thus, in the first iteration $r1$

Algorithm 1. `findMatchingRules`

```

Data: rules, matchingIntervals, i, N, matchingRules
1 if  $i < N$  then
2    $sortedEndpoints = rules.getSortedRulesEndpoints(i)$ 
3    $activeRules = []$ 
4   foreach  $endpoint \in sortedEndpoints$  do
5     if  $!activeRules.isEmpty()$  then
6        $matchingIntervals[i] = [lastEndpoint, endpoint]$ 
7        $findMatchingRules(activeRules, matchingIntervals, i + 1, N,$ 
8          $matchingRules)$ 
9     if  $endpoint.isLowerBound()$  then
10       $activeRules.add(endpoint.getRule())$ 
11    else
12       $activeRules.delete(endpoint.getRule())$ 
13     $lastEndpoint = endpoint$ 
14 else
15   if  $matchingRules.canBeMergedWith(\{activeRules, matchingIntervals\})$  then
16      $matchingRules.mergeWith(\{activeRules, matchingIntervals\})$ 
17   else
18      $matchingRules.add(\{activeRules, matchingIntervals\})$ 
19 return  $matchingRules$  /* contains sets of rules together with matching input
    intervals */

```

will be added to *activeRules*. Then, the current endpoint is saved as the *lastEndpoint* (line 11) and the next endpoint is processed by sweeping the line until the next endpoint is reached (cf. Fig. 2).

Now, if there is at least one active rule (line 5), then the interval in which these rules are active for that dimension is saved (line 6). Therefore, $matchingIntervals[i] = [0, 10)$.¹ Then, there is a recursive call to *findMatchingRules()*, to check the next dimension ($i + 1 = 1$), but only for the list of *activeRules* = [*r1*], and only in the *matchingIntervals* = $[[0, 10)]$. Therefore, the line being swept across across the second dimension only considers rules that lie within the matching intervals of the previous dimensions. This is visualized in Fig. 3.

At this point, $sortedEndpoints = [r1(0), r1(30)]$, so that in the first iteration of the for loop *r1* is added to the list of *activeRules* (line 9). In the second iteration, the list of intervals for which the active rules match is updated (line 6), such that $matchingIntervals = [[0, 10), [0, 30]]$. Then, the procedure is recursively called again. Since now $i = N = 2$, the else-part of the algorithm is executed (line 13). This means that the first matching rule is found

¹ The interval does not include the upper endpoint because for the value 10 also *r2* would be active.




Fig. 3. Line being swept through the second dimension in the interval $[0, 10)$

and since *matchingRules* is so far empty, it is added to the list (line 17): $matchingRules = [\{r1, [[0, 10), [0, 30]]\}$.

The algorithm then returns from the recursive call and deletes *r1* from *activeRules* because the current *endpoint* is $r1(30)$, such that it is now empty and all endpoints of the second dimension have been processed. Therefore,




Fig. 4. Line being swept through the second dimension in the interval $[10, 20)$

the algorithm returns to the first dimension, where the current *endpoint* is $r2(10)$ (cf. Fig. 2). This means that $r2(10)$ is added to *activeRules*, so that in the next loop iteration the line in Fig. 2 is swept until it hits the endpoint $r3(20)$. Hence, a recursive call is made where $activeRules = [r1, r2]$ and $matchingIntervals = [[10, 20)]$. This means that the line is swept across the second dimension in that interval (cf. Fig. 4), so that the algorithm will detect that

- $r1$ matches for the intervals $[[10, 20), [0, 10)]$,
- both $r1$ and $r2$ match for the intervals $[[10, 20), [10, 25)]$, and
- $r1$ also matches for the intervals $[[10, 20), (25, 30)]$.

Back in the first dimension, the current *endpoint* is $r3(20)$, which is added to *activeRules*, and the line is swept until it hits the closing endpoint $r1(30)$. Therefore, the line is swept across the second dimension in the interval $[20, 30]$. The current state of all *matchingRules* would then be visualized as in Fig. 5a.

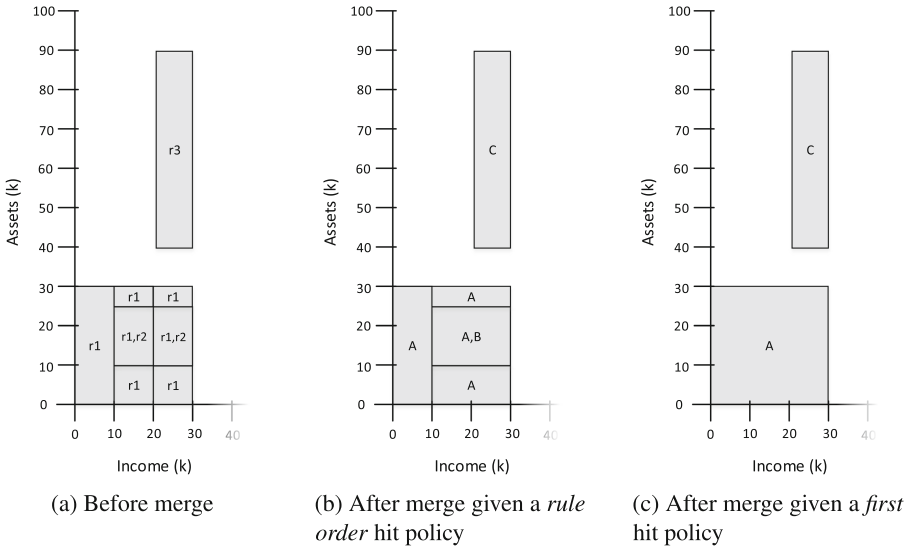


Fig. 5. Set of *matchingRules* found so far before and after merge

In this figure the rectangles are labeled with rules that match for this rectangle. For example, we found out that rules 1 and 2 both match for the rectangle $[[10, 20), [10, 25)]$, such that it is labeled $r1, r2$.

Based on this representation, it becomes obvious that there is room for improvement, because there are rectangles that are actually very similar, in the sense that they only differ in one dimension, but in this dimension they are adjacent. For instance, the two rectangles that are labeled with $r1, r2$ are the same with respect to the *Assets* dimension, i.e., they both cover the interval

[10, 25]. In the *Income* dimension they differ, but they are adjacent, because one covers the interval [10, 20) and the other one covers the interval [20, 30]. Thus, it seems that there is the possibility of merging pairs of rectangles into a single rectangle.

In general, two hyperrectangles with n dimensions can be merged if they overlap in $n - 1$ dimensions and are adjacent in the remaining dimension, and if the corresponding rules have the same output. This observation has already been described in [6]. However, [6] only deals with unique-hit tables. This means that a hyperrectangle only represents one rule, such that the corresponding output is unambiguous. In our case, however, more than one rule can belong to the same rectangle, so that the output of that rectangle depends on the hit policy of the underlying decision table.

For example, since the table in Fig. 1 has a *rule order* hit policy, the rectangles in Fig. 5a that are labeled $r1, r2$ will produce the output $[A, B]$, so that they can be merged with each other, but not with any of the remaining rectangles. The result of all possible merges given a *rule order* hit policy is illustrated in Fig. 5b. Now assume that the table has a *first* hit policy instead. In this case, all of the rectangles in the lower left corner in Fig. 5a will produce the same output (A). Therefore, the result after merging would be as shown in Fig. 5c.

Regarding our algorithm this means that whenever a new rectangle is found, it first checks if it can be merged with other rectangles found so far (line 14). This includes determining the output that is produced by the newly found rectangle based on the hit policy of the table. If merges are possible, the newly found and the existing rectangles are merged and then the combined result is put into the list of *matchingRules* instead of its parts (line 17). Therefore, after the line has been swept across the second dimension in the interval [20, 30], the list of *matchingRules* will actually look as visualized in Fig. 5b, because three rectangle pairs can be merged.

Eventually, the algorithm will have analyzed all *sortedEndpoints*, and the final list of *matchingRules* will have 20 elements in total, the last one being $\{[r4, r5], [(95, inf), [85, inf)]\}$. This list is visualized in Fig. 6. Given the visualization of this list, it becomes clear that we can now construct a new table that exactly contains the rules represented as rectangles in the figure. This table is shown in Fig. 7. It has 15 exclusive rules and is therefore a unique hit table as indicated by the letter U in the upper left corner. It has the same input-output behavior as the table in Fig. 1. However, due to the fact that all of its rules are exclusive it is more easily comprehensible and the set of possible outputs of the table is immediately visible. Furthermore, this representation of the table is more suitable for certain analysis tasks as described in Sect. 2.

4 Evaluation

We implemented our algorithm to disambiguate DMN decision tables in *dmn-js*, a DMN decision table editor developed by Camunda. Parts of our implementation reuse code written for the DMN decision table verification algorithms




Fig. 6. Decision table as a set of hyperrectangles

U	Income (k) <i>Number ≥ 0</i>	Assets (k) <i>Number ≥ 0</i>	Credit Rating <i>A,B,C,D,E,{A,B},{C,D},{C,E},{D,E}, {C,D,E}</i>
1	< 10	≤ 30	A
2	∈ [10..30]	< 10	A
3	∈ [10..30]	∈ [10..25]	A,B
4	∈ [10..30]	∈ [25..30]	A
5	∈ [20..40]	∈ [40..90]	C
6	∈ [30..60]	∈ [10..25]	B
7	∈ [40..80]	∈ [40..85]	C
8	∈ [40..80]	∈ [85..90]	C,E
9	∈ [40..80]	> 90	E
10	∈ [80..95]	< 40	D
11	∈ [80..95]	∈ [40..85]	C,D
12	∈ [80..95]	∈ [85..90]	C,D,E
13	∈ [80..95]	> 90	D,E
14	> 95	< 85	D
15	> 95	≥ 85	D,E

Fig. 7. Disambiguated table derived from the table in Fig. 1

described in [5,6]. Our evaluation is based on a total of 1000 synthetic decision tables that we generated for this purpose.

Specifically, we created tens sets of 100 decision tables. Each set contains tables with all combinations of {5, 10, 15, . . . , 50} rows and {3, 6, 9, . . . , 30} columns. The input columns are of type integer and the output column of type string. The rules are randomly generated by generating a random condition for each input. Each condition is made up of one of the operators in {<, >, ≤, ≥, =} and an integer, both chosen at random. The output of each rule is given by the rule number (as a string). Furthermore, each table has a multi-hit policy. In this

way, we generated tables that have a random number of overlapping rules and can therefore be disambiguated by our algorithm.

We applied the algorithm to all the tables in all ten sets and then averaged the execution times over the ten sets. The results are shown in Fig. 8 as a three-dimensional surface plot. The first observation is that the execution times depend more heavily on the number of rows than the number of columns. This is because the more rules there are the more intervals need to be checked for matching rules. However, the maximum average execution time was actually reached for the 50×6 -tables with 151 s; and also other tables with similar numbers of rules and columns lead to comparably high execution times.

The reason for that is the high number of overlapping rules of those tables. More specifically, tables with a high number of rules but a low number of columns have much more rules that overlap across all dimensions. This leads to a high number of *activeRules* in Algorithm 1 that need to be compared to all the other rules for potential merges (line 14). The average total number of overlaps of the rules of the tables in our data set is shown in Fig. 9. Here, the maximum number of overlaps is 1138 for the 50×6 -tables, which coincides with the maximum in Fig. 8.

Note that the number of overlaps for the tables with a high number of columns is nearly always zero. This is because the higher the number of columns, the harder it is for multiple rules to overlap in *all* of those dimensions. This shows that our disambiguation approach is especially suited for tables with a high number of rules and a lower number of columns.




Fig. 8. Execution times in seconds for multi-hit tables with up to 50 rows and 30 columns




Fig. 9. Total number of overlaps found in the tables

5 Related Work

Transforming decision tables to equivalent ones that are disambiguated, is highly related to the topic of normalization of decision tables, which has been brought up in [9] and re-used in [10]. Normalization of decision tables is based on normalization rules in relational databased design and also aims at increasing the understandability of the tables. To this end, [9,10] describe three normal forms of decision tables.

The first form is mainly concerned with the layout of the table and the structure of its logical expressions. For example, rules of a decision should be represented as rows in the table (DMN also allows a rules-as-columns layout). Moreover, the individual input conditions must always be connected conjunctively, never disjunctively. Second normal form requires that, unless an input is explicitly marked as irrelevant for a particular rule, all conditions of a rule must be relevant for the output of that rule. For example, in rule 4 in the table in Fig. 1, the *Assets* input is irrelevant and therefore marked with “-”. Finally, third normal form demands that the input of the tables should be independent from each other. Thus, it should not be possible to infer the value of an input condition from the values of the other input conditions.

Conforming to these normal forms increases the readability and understandability of the table. However, they assume that the table already consists of exclusive rules, i.e., conforms to the *unique*-hit policy. Hence, our approach starts one step earlier by first disambiguating tables with overlapping rules into tables with exclusive rules. Subsequently, the normalization rules of [9,10] can be applied to further improve the tables.

As already described in Sect. 2 the algorithm for table disambiguation is based on algorithms for detecting overlapping and missing rules as well as simplifying tables by merging similar rules [5,6]. The table simplification algorithm

assumes a *unique*-hit table. Our algorithm derives a *unique*-hit table from a table with overlapping rules that is inherently simplified because it applies rule merging on the fly, i.e., during the identification of exclusive rules, as described in Sect. 3.

6 Conclusion

In this paper, we presented an algorithm for the disambiguation of DMN decision tables. The algorithm has been implemented and evaluated against a set of synthetic decision tables and we could illustrate interesting characteristics about it such as the influence of the number of overlapping rules. In fact, it is not obvious how many overlapping rules an arbitrary DMN decision tables can actually have. Our evaluation showed that for tables with a high number of rules and a low number of columns, the number of overlaps can get very high. But also the domains of the input variables should be a factor, since, for example, integers have much more possibilities for overlaps than booleans.

Therefore, in future work we would like to find out more about the maximum number of overlapping rules given a DMN table with a certain number of rows and columns and where the columns have certain domains. This will provide a way to determine for which kinds of tables the disambiguation procedure is most suitable and will produce more understandable and less error-prone tables.

References

1. Batoulis, K., Haarmann, S., Weske, M.: Various notions of soundness for decision-aware business processes. In: Mayr, H.C., Guizzardi, G., Ma, H., Pastor, O. (eds.) ER 2017. LNCS, vol. 10650, pp. 403–418. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-69904-2_31
2. Batoulis, K., Weske, M.: Soundness of decision-aware business processes. In: Carmona, J., Engels, G., Kumar, A. (eds.) BPM 2017. LNBIP, vol. 297, pp. 106–124. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-65015-9_7
3. Bazhenova, E., Haarmann, S., Ihde, S., Solti, A., Weske, M.: Discovery of fuzzy DMN decision models from event logs. In: Dubois, E., Pohl, K. (eds.) CAiSE 2017. LNCS, vol. 10253, pp. 629–647. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-59536-8_39
4. Bentley, J.L., Ottmann, T.A.: Algorithms for reporting and counting geometric intersections. *IEEE Trans. Comput.* **28**(9), 643–647 (1979)
5. Calvanese, D., Dumas, M., Laurson, Ü., Maggi, F.M., Montali, M., Teinmaa, I.: Semantics and analysis of DMN decision tables. In: La Rosa, M., Loos, P., Pastor, O. (eds.) BPM 2016. LNCS, vol. 9850, pp. 217–233. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-45348-4_13
6. Calvanese, D., Dumas, M., Laurson, Ü., Maggi, F.M., Montali, M., Teinmaa, I.: Semantics, analysis and simplification of dmn decision tables. *Information Systems* (2018)
7. Hasić, F., Smedt, J.D., Vanthienen, J.: Augmenting processes with decision intelligence: principles for integrated modelling. *Decis. Support Syst.* **107**, 1–12 (2017)

8. OMG: Decision Model and Notation, Version 1.1, May 2016
9. Vanthienen, J., Snoeck, M.: Knowledge factoring using normalization theory. In: International Symposium on the Management of Industrial and Corporate Knowledge (ISMICK), pp. 97–106 (1993)
10. Von Halle, B., Goldberg, L.: The Decision Model: A Business Logic Framework Linking Business and Technology. Taylor and Francis Group, Boca Raton (2010)

Smart Infrastructures



Using Blockchain Technology for Business Processes in Purchasing – Concept and Case Study-Based Evidence

Stefan Tönnissen^(✉) and Frank Teuteberg

Universität Osnabrück, Osnabrück, Germany
{stoennissen, frank.teuteberg}@uni-osnabrueck.de

Abstract. How can blockchain efficiently and effectively support purchasing processes? This paper addresses this question on the basis of a case study and an analysis of expert interviews. Furthermore, a qualitative content analysis highlights the challenges, barriers and perceived benefits associated with the blockchain technology. The paper concludes that blockchain technology, with its ability to connect to existing ERP systems, has the potential to make processes in a company's purchasing environment more efficiently and transparent by the use of smart contracts on the blockchain.

Keywords: Blockchain · Purchasing · Supply chain · Smart contracts

1 Introduction

At the 25th International Purchasing and Supply Education and Research Association Conference 2016 in Dortmund, the research question “How can blockchain technology improve process flows and transparency between buyers and suppliers?” was explored in a research framework [1]. One study found that the current use of modern IT tools, especially in purchasing and related supply chain management, is not very advanced [2]. Obviously, companies are not well prepared for the requirements of networking through Industry 4.0 [2]. The purpose of this article is to examine how the use of blockchain technology can affect the existing processes and systems of purchasing raw materials. We focus on the procurement of raw materials, since we want to incorporate the requirements of the networking raw material process machines in Industry 4.0. In our concept, we also focus on Enterprise Resource Planning systems (ERPs) and a holistic view of the procurement process. Thus our research guiding question is:

How can business processes in purchasing be designed more efficiently and effectively using blockchain technology?

The article is structured as follows. After the introduction, Sect. 2 explains the theoretical background of blockchain technology as well as purchasing in the supply chain concept. In addition, the methodological approach of the article as well as related work is described. We then use interviews in Sect. 3 to analyze the current issues and challenges, the perceived benefits, and barriers associated with the blockchain technology. To do

this, we classify our results using an empirical-to-conceptual approach to obtain a more accurate picture. The following is the explanation of the case study and development of a process flow with ERPs and blockchain technology as an artifact in Sect. 4. The discussion of the results will take place in the following Sect. 5. The results of the experts' evaluation of the artifact will be presented in Sect. 6. The article ends with a summary and an outlook in Sect. 7.

2 Theoretical Background

2.1 Blockchain Technology and Smart Contracts

The blockchain technology was already described in a white paper by Satoshi Nakamoto in 2008 and has been in use since 2009 with the crypto-currency Bitcoin [3]. A blockchain is a stringing (concatenation) of data, which are combined into individual blocks and stored on all users' computers. This succession of data into blocks results in a sequence that reflects the course of transactions like a chain. All the data blocks are protected against subsequent changes by means of cryptographic methods, so that over time a gapless chain of linked data blocks is created. The inclusion of a new record in the blockchain requires the passage of a so-called consensus mechanism that runs across the network of all participants and is used to reach an agreement among all participants in the blockchain network about the correct state of data on the blockchain. This ensures that the data is the same on all nodes in the network [4]. The best-known consensus mechanism is proof-of-work, in which the computer has to execute a complicated mathematical algorithm with great effort. Only after successful execution a new data block can be generated in the blockchain, which must be checked by the other computers in the peer-to-peer network before inclusion in the blockchain [5]. In addition to the data, each block contains a timestamp as well as the hash value of the previous block. The blocks are protected against subsequent changes by means of cryptographic methods, so that a continuous chain of linked data blocks is formed over time [5].

2.2 The Procurement Function in Supply Chain Management

The procurement of goods and services is an essential part of commercial enterprises. The importance of purchasing goods and services is evident in that, for example, car manufacturers often spend over 50% of their sales on procurement [6]. Within Porter's value chain, procurement as a secondary activity takes a cross-cutting role in primary activities [7]. The procurement process is an essential element of supply chain management that describes the cross-company coordination of material and information flows [31]. Procurement as a component of supply chain management is in the focus of digitization and networking, since this area acts as an interface between the company's internal units and the procurement market or supplier [2].

2.3 Related Works

Based on a previous review of the literature, we were able to find some related contributions. Kshetri [9] examines the impact of blockchain on key supply chain management goals and concludes that there is a high potential for achieving the goals. Biswas et al. [10] address the need for traceability in the wine supply chain and propose a blockchain-based system with the result that a high quality information management system could solve the problems of the wine industry. Korpela et al. [11] attack the blockchain for business process integration and explain how integration can succeed. Hackius and Petersen [12] use a survey of logistics professionals to show the potential of the blockchain for logistics. Compared to these rather conceptual contributions in literature to the possibilities of the blockchain in the supply chain management we will present a case study and an analysis of barriers, challenges and benefits of blockchain technology based on conditions of a real company.

2.4 Methodical Approach

To answer our research question, we first conduct a qualitative content analysis of interviews. With qualitative content analysis it is possible to classify words into content categories [8]. Based on the classification of current problems and challenges as well as advantages and obstacles, we have developed a case study. This case study addresses the realities of a business and attempts to solve previously identified problems and challenges using blockchain technology. The evaluation is then carried out by interviewing experts through a standardized questionnaire.

3 Analysis of Interviews

We first conducted an analysis of interviews with blockchain experts on the opportunities and challenges for blockchain technology for logistics processes/supply chain management. For this purpose, we entered the search string “blockchain” and “interview” and “supply chain” or “logistic *” “in Google for the period 01.01.2017 to 01.31.2018 and received 35.400 results. Based on our assumption that the results of the first pages reflect the relevance of Google’s search algorithms [13], we used the titles and short texts to analyze the results to filter out the interviews relevant to our research question. We first transferred the first 20 interviews in an excel file with the title fields, who was interviewed, in which role is the interviewee, who interviewed, when did the interview take place, on which source and when was the interview found. The interviews were then evaluated in a next step regarding the following questions:

- What problems or challenges are seen in logistics or in supply chain management?
- What are the advantages of using the blockchain technology in the logistics industry?
- What are the obstacles to using blockchain technology?

Due to the fact that not all interviews were able to provide the necessary information to answer the questions, further interviews from the Google results list were gradually taken over into the Excel file and examined with regard to the questions.

In the end, 35 interviews were evaluated (see also <https://tinyurl.com/ybbuw7fd>). The results for the challenges, benefits, and obstacles are further illustrated by an empirical-conceptual approach according to Nickerson et al. [14] and have been classified with the results in Table 1 (see also <https://tinyurl.com/ybbuw7fd>).

Table 1. Results of the interviews with classification.

Problems/Challenges		Advantages		Obstacles	
Class	Number of responses	Class	Number of responses	Class	Number of responses
Process	9	Process	36	Adaptation	7
Trust	5	Transparency	28	Trust	6
Conditions	4	Fraud	8	Technology	4
Data	2	Costs	5	Organization	4
IT Security	2	Organization	4	Business	2
Fraud	1	IT security	4	Legal	2
Costs	1	Collaboration	3	Financing	1
Standards	1	Trust	3	Operating cost	1
		Risk	2	Network effect	1

The class process contains both the external view with the cooperation with business partners in a logistics chain as well as the internal view for the integration of the primary and secondary activities in a company. A key issue of the current processes was mentioned by the interviewees in terms of paper-based operations (e.g. shipping documents, customs papers, export declarations, warehouses). Due to the high relevance of the class process both in terms of problems/challenges and the perceived benefits of the blockchain technology for the processes as well as the class transparency we subdivided them further into subclasses. According to Becker and Kahn [15], a process is “the content-related, temporal and logical sequence of activities that are necessary for processing a business-relevant object.” The subclasses for the class process and for the class transparency are based on the empirical conceptual approach Nickerson et al. [14] have developed. In the class process, the subclass time plays an important role according to the evaluation of the interviews. The perceived benefits of using blockchain technology focus on timely processing, exemplary comments are “... to have access to the right information at the right time” [16]. Furthermore, an improvement in the quality of the processes is expected, such as e.g. “...better tracking of orders, reducing errors and better fraud detection” [17]. The avoidance of fraud by transaction processing in real time and the immutability of the data on the blockchain is another supposed advantage as well as the potential of automated process processing by smart contracts on the blockchain. Finally, the interviewees emphasized the importance of security in the processes by “... logistics industry wants to see improved connectivity, efficiency and security thanks to blockchain” [18]. The class transparency contains an external and an internal view. The external view of transparency refers to the exchange of information with business partners within a logistical process in a logistical chain [19]. The internal

view in this context means both an insight into the own order status of a logistical process as well as an insight into the entire logistical chain with the possibility of backward and forward traceability of an order. Blockchain offers interviewees the advantage of being able to provide proof of possession as well as proof of transport [20]. Closely associated with this is the benefit of status tracking of the status of the flow of goods, highlighted by e.g. “Logistics service providers, for example, can document all incidents along the supply chain completely, unchangeable and visible to everyone” [21]. At the same time, documentation of a history of the logistical process, highlighted e.g. “With the distributed database, network participants can directly engage in transactions and see the history of all transactions” [22]. Parnell [32] emphasizes the importance of real-time processing with “real-time sharing of information about process improvements and maintenance”. The immutability of the data in a blockchain is emphasized by “The information captured in each transaction is agreed upon by all members of the business network; once there is a consensus, it becomes a permanent record that cannot be changed” [32]. In the class obstacles the interviewees see a significant aspect in the difficulty of adapting a blockchain-based application. Exemplary comments on this are “Customer engagement is about creating information systems that are truly accepted. This is less a matter of technology and more a matter of having an approach that inspires all parties involved, not just customers to work in incremental steps with tangible results towards a solution.” [16]. Closely related to the difficulty of adaptation is the lack of confidence that exemplifies “... and adopting a new mindset around a decentralized network with no central control” [32].

4 Case Study

To answer the research question, we conduct a case study that, according to Ridder [24], has the advantage of a more detailed description and more detailed analysis. On this basis, the questions about the “how” and “why” can be answered more easily [24]. Our case study is suitable in our research subject because a current phenomenon (blockchain in purchasing) in a real and practical context (the company Schmitz Cargobull AG) is examined [25]. According to Brüsemeister [26], a case-by-case study is also useful if it provides access to a hitherto little-explored social area, which points to the use of blockchain technology in the integration of “machine-to-ERP-to-blockchain-to-ERP-to-machine” applications.

The company in our case study is Europe’s leading manufacturer of semi-trailers and trailers for temperature-controlled freight, general cargo and bulk goods with an annual production of around 58,000 vehicles and around 5,700 employees [27]. Within the value chain of Schmitz Cargobull AG, the procurement of raw materials, primary products, consumables and tools occupies a significant position. A logistics manager explained in one of our interviews that one goal of supply chain management is the optimization of transparency about capacities in the production and procurement network [23]. Since 2002, the entire order processing has been carried out via the AXIT logistics platform AX4. Through the connection of the own SAP system as well as all suppliers the data exchange in a procurement process can be automated [28]. In the

following, a process model will be constructed based on existing information, taking into account both the challenges of analyzing the interviews and the technical capabilities of the blockchain technology. Furthermore, we take the Industry 4.0 concept into account by linking the industrial infrastructure such as machines in the concept.

In our case study, a networked machine determines a material bottleneck and then automatically makes a demand request to the ERPs. In the ERPs, a planning run is carried out on the basis of the existing plant stocks, and a procurement proposal is created based on the resulting requirements calculation. This procurement proposal is written to the blockchain as a purchase requisition with the characteristics material number, material name, quantity, unit of measure, delivery date, recipient via the blockchain gateway of the ERPs. The suppliers have installed a smart contract on the blockchain, which recognizes the requirement and transfers the data record via the blockchain gateway to the supplier's ERPs as a demand request. At the same time, after submitting the data to the ERPs, the smart contract sets a flag on the blockchain that enables the customer to identify the potential supplier's activity. The supplier can now use the information (material, quantity, delivery time, location, quality, etc.) to create a quotation and write this as a record via the blockchain gateway to the blockchain. The customer has installed a smart Contract on the blockchain, which recognizes the suppliers' offers on the blockchain and checks them independently using stored rules (price, quality, supplier evaluation, etc.) and transmits the data via the blockchain Gateway to the ERPs. In the ERPs, the incoming quotations are subjected to an offer check according to defined criteria such as quantity, price, delivery date, quality, supplier evaluation, etc. The cheapest offer is then written to the blockchain via the blockchain gateway. The smart Contract then creates an order for the best supplier and confirms the offer. Thus, according to contract law a valid purchase contract came about [30]. The job is also written to the blockchain. The supplier has a smart contract running, which recognizes the order, checks it and, after transfer to the ERPs, includes it in the capacity planning of production for these raw materials. The scheduling for the production order is written to the blockchain via the blockchain gateway and transferred directly to the customer's ERPs via a smart contract. This receives a status of the manufacturing process. After completion of the manufacturing process, the status is first rewritten to the blockchain and then the material is delivered to the customer. The delivery of the raw materials is written to the blockchain with quantity, quantity unit, material number, quality, etc. and transferred to the customer's ERPs via a smart contract. The customer then only has to confirm the proposed goods receipt in the ERPs during the physical goods receipt. Then the goods receipt in the ERPs is written to the blockchain via the blockchain gateway and the supplier receives the status of the goods receipt at the customer.

Due to the high importance of the processes, a detailed explanation of the individual process steps with the corresponding systems and transactions is presented in the following process flow. First, the entire process is divided into five individual sub-processes. On the customer side, the first step includes requirement recognition and, via the requirements requisition and purchase requisition, leads to the supplier side with the data transfer and processing. In the subsequent second step, the supplier's offer leads to an examination of the offer by the customer towards an order. This is transferred in the third step in the capacity planning and scheduling at the supplier to the finished

notification of the manufacturing process to the customer via a status message. The fourth step prepares the delivery to the customer and writes the data of the goods delivery to the blockchain, which are taken over by the customer directly into the ERPs and form the basis for the posting of the physical goods receipt. This goods receipt at the customer finds in the fifth step and end with the corresponding status message to the supplier (Fig. 1).

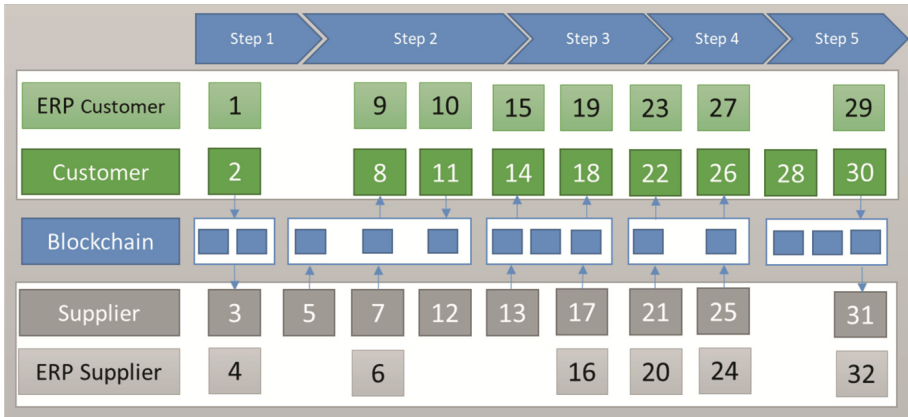


Fig. 1. Process flow with ERP systems and blockchain.

Step 1: (1) Demand request, (2) Purchase requisition, (3) Data transfer ERP, (4) Demand request, (5) Activity indicator. **Step 2:** (6) offer, (7) data submission offers to blockchain (8) takeover offer, (9) offer review, (10) determine favorable offer, (11) order on the blockchain, (12) smart contract checks order, (13) order confirmation, (14) smart contract checks order confirmation, (15) take over status in ERP. **Step 3:** (16) scheduling, production order, (17) appointment on blockchain, (18) smart contract recognizes appointment, (19) take over appointment, (20) finish, (21) write status in blockchain, (22) smart contract recognizes status and takes over in ERP, (23) takeover status. **Step 4:** (24) Determine delivery data, (25) Write delivery data in blockchain, (26) smart contract recognizes delivery data and takes over in ERP, (27) Transfer of goods receipt data, (28) Physical goods receipt. **Step 5:** (29) Posting Goods Receipt, (30) blockchain Indicator, (31) smart contract recognizes status Goods Receipt, (32) Transfer to ERP, and Basis for Invoicing.

The blockchain will continue to be used as middleware for exchanging status-related information between the customer's ERPs and suppliers. Thus, the transparency in the process of procurement of raw materials for both the buyer and the supplier is significantly increased. The status-relevant information is written to the blockchain in real time and is available in a timely manner (Tables 2 and 3).

Table 2. Status tracking on the blockchain.

Status points	Explanations
(5) Activity indicator	The customer uses the indicator to identify whether a supplier has accepted the purchase requisition for processing
(17) Appointment on blockchain	The customer receives the scheduling of the production order from the supplier
(21) Status of production	Completion of the manufacturing process at the supplier is transmitted to the customer
(25) Delivery dates	The data about the delivery of the goods will be published
(30) Goods receipt	The supplier receives the information with the indicator that the goods receipt has been made with the customer

70% of Industry 4.0 strategies are aimed at increasing productivity and efficiency [2]. To improve efficiency in our concept, we set up the subsequent smart contracts on the blockchain in the processes.

Table 3. Functions of smart contracts.

Smart contracts	Functions
(3)	The smart contract recognizes the purchase requisition relevant to the supplier and transfers the data to the gateway to the ERPs
(8)	The smart contract recognizes the purchase requisition offer and transfers the data to the gateway to the ERPs
(12)	The order from the customer is matched by a smart contract with the previously submitted offer of the supplier
(14)	The smart contract checks the order confirmation and performs an adjustment to the order
(18)	The smart contract recognizes the supplier’s appointment for the procurement process and transmits the date via the gateway to the ERPs
(22)	The smart contract recognizes the status of production set by the supplier and transmits the status via the gateway to the ERPs
(26)	The supplier’s delivery data is taken over by a smart contract and transferred via the gateway to the customer’s ERPs
(31)	The customer writes the status of the goods receipt to the blockchain, this status is recorded by a smart contract and transferred via a gateway to the ERPs of the supplier

5 Discussion of Case Study Results

The result of our case study shows a process that is characterized by the consideration of existing ERPs both at the manufacturer and the supplier. The blockchain is switched as a middleware between these two ERPs and thus provides the participants with a consistent and transparent database with the guarantee, the immutability and traceability of the data of the logistical processes in purchasing. With the blockchain as a common database, media breaks will be avoided in the future. The connection of the existing

ERPs is done by blockchain gateways, which are set up on the part of the operators of the ERPs. Thus, the sovereignty over the data import as well as data export remains in the hands of the operator. With this conception, the adaptation of the blockchain into the operational processes could succeed and the hurdle of the initially low confidence could be overcome. The connection of further suppliers could succeed on this basis without major obstacles, and thus lay the basis for further acceptance of the blockchain. Our case study shows significant improvements in the process, since the number of interactions by a person could be significantly reduced from the needs assessment through a networked machine at the manufacturer to the receipt of the required raw materials. Thus, both the time required for the processing of the processes can be reduced and the process reliability and quality can be increased. The use of smart contracts contributes significantly to the automation of process steps and thus ensures greater security in process processing. All companies involved in the process receive transparency about the overall process via the status reports in real-time on the blockchain, and can purposefully control their material supply on the basis of this.

6 Expert Interviews

For the evaluation of our concept, we consulted experts from the operational practice and asked to fill out a standardized questionnaire (questionnaire and profile of the participants see also <https://tinyurl.com/ybbuw7fd>). The experts participated in a workshop on blockchain technology on November 11th 2017. We invited 11 participants for the survey. The standardized questionnaire was made available to the participants on November 12th 2017. The respondents gave the following answers to the questions. For the question “How realistic do you think the process model with ERP and blockchain is?”, 60% of respondents think it is a realistic use case. 30% of respondents are rather critical or uncertain about the assessment. The realistic estimate previously given with 60% can be found in comments such as “Technical feasibility I consider realistic” or “This process happens in all companies in the world every day”. In addition, there are critical voices such as “However, there will be distrust in the technique that probably outweighs” or “The system also depends on the number of companies that work with it. I cannot imagine a widespread distribution in the next 10 years”. The question “What is your confidence in the blockchain technology?” was answered by 50% for rating 3. In our odd rating scale, rating 3 represents a neutral middle category. The high number of responses could therefore suggest an uncertain judgment of the participants, since the middle scale point can also be used as a flight category [33]. The comments of the participants confirm the suspicion of the uncertain judgment. Exemplary comments are “The confidence is neither particularly high nor particularly deep, because I could not gather any experience with it” or “Due to my low level of knowledge about this technology, I find it difficult to make a statement about a manipulation of the individual steps”. Answering the question “Do you think that blockchain technology has reached a necessary maturity level for a” real “mission?” led to a balanced result as 50% of respondents answered yes and 50% answered no. The proponents commented on their decisions, among others with “I think the technology can already be used in a limited

field of application for special applications” or “Today almost all production facilities are equipped with the appropriate interfaces to introduce these systems nationwide”. The interlocutors justify their opinion with “I think this step to the introduction is currently too big. It would have first created general trust in the technology” and “because I cannot assess how far the degree of maturity has progressed”. In our standardized questionnaire, we asked the participants for an assessment as to whether an effective purchase agreement pursuant to contract law had been concluded. The result was 80%, no. The reasons for this assessment are exemplary “The contract law is in my opinion not for machines or IT systems”, “Because the behavior is too digitalized” or “offer and acceptance cannot be decided by a digital form”. The participants, who claim that an effective sales contract was made with 20%, justified this with “if on eBay the automatic bidder system was activated by the buyer, then also a sales contract came about” or “one could assume that a purchase contract has been concluded by offer and acceptance”.

7 Summary and Outlook

We can summarize that the processes of purchasing using blockchain technology have a potential that can solve today’s problems and challenges. The blockchain technology has the potential to meet the expectations of the interviews with regard to improving processes and increasing transparency. By using the smart contracts, numerous process steps could be automated and efficiency gains achieved through real-time processing of blockchain. By integrating networked machines, the integration of Industry 4.0 concepts succeeds. The validation of our concept by a survey of experts from different companies has shown that the predicted use case with the integration of a blockchain in the existing ERPs of customers and suppliers is realistic. However, interviewees also found barriers to successful adaptation of such a solution. It is unclear how adaptation could take place via a global logistics chain. The trust required between the business partners was interpreted by the interviewees as a clear barrier to integration.

The case study is unrepresentative due to the company’s choice, but demonstrates a strong ability to collaborate with business partners based on a high prevalence of ERPs in manufacturing companies. In addition, existing IT concepts for connecting ERPs across company boundaries could also be able to solve the challenges of purchasing. Blockchain technology competes with electronic data interchange systems such as EDI, EDIFACT, AXIT, etc. This entry could help to discuss whether existing IT systems for enterprise networking or business-to-business data exchange are inferior to blockchain technology. Here further investigations regarding a cost, benefit and risk assessment are necessary.

The dissemination of networking of machines with ERPs and subsequently with business partners is facilitated by standardizing both the information technologies required for this purpose, e.g. gateways for the connection to the blockchain as well as data structures and the document structures relevant for a global supply chain process favors. These requirements for global standardization cannot be met to this day [11]. This makes it difficult to adapt a corresponding blockchain-based application to existing

system landscapes of companies. In addition, a supplier is wondering why he should invest in a technology that has not yet surpassed pilot status in some applications. There are also questions regarding the performance, scalability or operational readiness of blockchain-based applications. Another key aspect of successful adoption is the confidence of potential users in both the technology [29], which is resource intensive and requires storage space for the complete blockchain [11], as well as in a business model that works without a central and trustworthy instance [22]. How can it be ensured that all relevant business partners voluntarily entrust their data to the blockchain [22]?

References

1. Foerstl, K., Schleper, M.C., Henke, M.: Purchasing and supply management: from efficiency to effectiveness in an integrated supply chain. *J. Purchasing Supply Manag.* **23**(2017), 223–228 (2017)
2. The purchasing and logistics working group of the Schmalenbach-Gesellschaft für Betriebswirtschaft e.V. (2017): Digitalisierung und Vernetzung in Einkauf und Supply Chain Management. In: Krause, S., Pellens, B. (Hrsg.): Betriebswirtschaftliche Implikationen der digitalen Transformation. *ZfbF-Sonderheft 72/17*
3. Yli-Huumo, J., Ko, D., Choi, S., Park, S., Smolander, K.: Where is current research on blockchain technology? - A systematic review. *PLoS ONE* **11**(10), e0163477 (2016)
4. Swan, M.: *Blockchain, Blueprint for a New Economy*. O'Reilly, Farnham (2015)
5. Holotiuk, F., Pisani, F.; Moormann, J.: The impact of blockchain technology on business models in the payments industry. In: Leimeister, J.M., Brenner, W. (Hrsg.): *Proceedings der 13. Internationalen Tagung Wirtschaftsinformatik (WI 2017)*, St. Gallen (2017)
6. Bungard, P.: *CSR und Geschäftsmodelle, Auf dem Weg zum zeitgemäßen Wirtschaften*. Springer, Heidelberg (2018). <https://doi.org/10.1007/978-3-662-52882-2>
7. Porter, M.: *Competitive Advantage: Creating and Sustaining Superior Performance*. The Free Press, New York (1985)
8. Elo, S., Kyngäs, H.: The qualitative content analysis process. *J. Adv. Nurs.* **62**(1), 107–115 (2007)
9. Kshetri, N.: Blockchain's roles in meeting key supply chain management objectives. *Int. J. Inf. Manag.* **39**(2018), 80–89 (2018)
10. Biswas, K., Muthukumarasamy, V., Tan, W.L.: Blockchain based wine supply chain traceability system. In: *Future Technologies Conference (FTC)* (2017)
11. Korpela, K., Hallikas, J., Dahlberg, T.: Digital supply chain transformation toward blockchain integration. In: *Proceedings of the 50th HICSS* (2017)
12. Hackius, N., Petersen, M.: Blockchain in logistics and supply chain: trick or treat? In: *Proceedings of the Hamburg International Conference of Logistics (HICL)* (2017)
13. Google Inc.: How search works. <http://www.google.com/intl/ALL/search/howsearchworks/>. Accessed 02 Feb 2018
14. Nickerson, R.C., Varshney, U., Muntermann, J.: A method for taxonomy development and its application in information systems. *Eur. J. Inf. Syst.* **22**(3), 336–359 (2013)
15. Becker, J., Kahn, D.: *Der Prozess im Fokus*. In: Becker, J., Kugeler, M., Rosemann, M. (Hrsg.) (2012): *Prozessmanagement. Ein Leitfaden zur prozessorientierten Organisationsgestaltung*, 7. Auflage. SpringerGabler, Heidelberg (2012). 10.1007/978-3-642-33844-1_1
16. logifint (2017): Blockchain & logistics: an interview with Louis De Bruin. <http://logisticsandfintech.com/blockchain-logistics-interview-louis-de-bruin>. Accessed 30 Jan 2018

17. Burns, J.: Dnata and its partners test the use of blockchain technology. www.aircargoweek.com/dnata-partners-test-blockchain-technology/. Accessed 31 Jan 2018
18. Henderson, J.: Blockchain technology “set to revolutionise” logistics industry (2017). www.supplychindigital.com/technology/blockchain-technology-set-revolutionise-logistics-industry. Accessed 31 Jan 2018
19. Lamming, R.C., Caldwell, N.D., Harrison, D.A., Phillips, W.: Transparency in supply relationships: concept and practice. *J. Supply Chain Manag.* **37**(4), 4–10 (2001)
20. Nördinger, S.: Darum sollten sich Industrieunternehmen mit Blockchain beschäftigen (2017). <https://www.produktion.de/specials/revolution-blockchain/darum-sollten-sich-industrieunternehmen-mit-blockchain-beschaeftigen-264.html>. Accessed 30 Jan 2018
21. Gläser, T.: Logbuch ohne zentrale Autorität (2017). <https://www.it-zoom.de/dv-dialog/e/logbuch-ohne-zentrale-autoritaet-18512/>. Accessed 30 Jan 2018
22. Poll, D.: Logistik: Vorteile durch Blockchain in der Supply Chain (2017). <https://www.produktion.de/technik/logistik/logistik-vorteile-durch-blockchain-in-der-supply-chain-115.html>. Accessed 02 Feb 2018
23. Schonefeld, P.: Planung im Kreislauf (2013). <https://beschaffung-aktuell.industrie.de/supply-chain-management/planung-im-kreislauf/>. Accessed 02 Feb 2018
24. Ridder, H.-G.: The theory contribution of case study research designs. *Bus. Res.* **10**(2), 281–305 (2017)
25. Yin, R.: *Case Study Research: Design and Methods*. Sage, Thousand Oaks (2002)
26. Brüsemeister, T.: *Qualitative Forschung. Ein Überblick*, Wiesbaden (2008)
27. Schmitz Cargobull AG: Geschäftsjahr 2016/2017. 10.000 Fahrzeuge mehr produziert als geplant. Marktführer Schmitz Cargobull steigert Umsatz auf mehr als zwei Milliarden Euro (2017). https://www.cargobull.com/de/Detail_news-523_213_383.html. Accessed 02 Feb 2018
28. AXIT: Schmitz Cargobull Success Story (2016). https://www.axit.de/images/successstories/PDFs/AX4_Success_Schmitz-Cargobull_d.pdf. Accessed 02 Feb 2018
29. Nördinger, S.: Darum passen Blockchain und Industrie 4.0 zusammen (2017). <https://www.produktion.de/specials/revolution-blockchain/darum-passen-blockchain-und-industrie-4-0-zusammen-278.html>. Accessed 30 Jan 2018
30. Raskin, M.: The law and the legality of smart contracts. *Geo. L. Tech. Rev.* **305** (2017) https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2959166. Accessed 30 Aug 2017
31. Scholz-Reiter, B., Jakobza, J.: *Supply Chain Management – Überblick und Konzeption*. HMD 207/1999 (1999)
32. Parnell, G.: How blockchain can strengthen supply chain links (2017). <https://theloadstar.co.uk/guest-blog-blockchain-can-strengthen-supply-chain-links>. Accessed 30 Jan 2018
33. Duller, C.: *Einführung in die Statistik mit Excel und SPSS. Ein anwendungsorientiertes Lehr- und Arbeitsbuch*. 3. Auflage. Springer Gabler (2013). <https://doi.org/10.1007/978-3-642-37859-1>



Developing a Multiple-Objective Demand Response Algorithm for the Residential Context

Dennis Behrens^(✉), Thorsten Schoormann, and Ralf Knackstedt

Department of Information Systems,
University of Hildesheim, Universitätsplatz 1, 31141 Hildesheim, Germany
dennis.behrens@uni-hildesheim.de

Abstract. Energy grids are facing various challenges, such as new appliances and volatile generation. As grid reliability and cost benefits are endangered, managing appliances becomes increasingly important. Demand Response (DR) is one possibility to contribute to this task by shifting and managing electrical loads. DR can address multiple objectives. However, current research lacks of algorithms addressing these objectives sufficiently. Thus, we aim to develop a DR algorithm that considers multiple DR objectives. For evaluation, we implemented the algorithm and formulated demonstration cases for a simulation. The evaluated algorithm contributes for example to users and energy providers by realizing various benefits.

Keywords: Demand Response · Demand side management
Algorithm engineering · Greedy heuristic · Optimization

1 Introduction

Due to technological improvements and changing environmental conditions, energy grids are facing various challenges, such as growing energy demand through an increasing population and altering consumption patterns through new living behaviors as well as a rising pollution [1]. As a result, decentralized and sustainable concepts like volatile energy (e.g., photovoltaic (PV), wind and water) and new appliances (e.g., Electric Vehicles (EVs)) are implemented [2]. These concepts result in a more volatile energy generation and consumption which affects supplier and demand side. Possible effects can be critical peaks, contingencies, a volatile load profile, and disadvantages in the market performance as well as an insufficient infrastructure usage (e.g. [3]). These might result in blackouts, brownouts, shortages, a high spinning reserve and can endanger the reliability of energy grids. Especially the residential context that consumes about one-third of the energy (e.g., [4, 5]) and faces major transformation (e.g., implementation of EVs and decentralized energy generation) can contribute to this problem. A potential answer for this is managing appliances (e.g., washing machines, dishwashers, EVs and lightning) (e.g., [6]).

One possibility to manage appliances is given by Demand Response (DR). DR focuses on optimizing consumption patterns (e.g., according to external (pricing) signals).

DR algorithms can be heterogeneous regarding, for example, objectives, optimization methods or communication structures [7]. Optimization in DR is carried out by shifting or managing loads, provoked by incentive-based programs [8] such as different types of dynamic pricing (DP) (e.g., time-of-use pricing (ToUP), critical-peak pricing (CPP), and real-time pricing (RTP) [3, 9]). To carry out the complex DR task with changing pricings, supporting infrastructure is needed (e.g., Smart Metering), that is not widespread. Hence, fast changing pricing schemes are not accomplishable and more static pricings are chosen.

A popular and more static DP scheme is ToUP with one low pricing interval (off-peak hours) and one high pricing interval (peak hours). The pricing signal therefore is known and predictable, and the communicational effort is low. However, this approach may result in misleading incentives (e.g., switching on all appliances at the same time, resulting in a new peak). A possible solution is a pricing function that extra charges peaks and a fluctuating load profile. Similar types of pricing signals already exist [10] (e.g., load depended pricing and CPP). However, current DR algorithms usually take only a single factor or objective into account (cf. [11, 12]) and we lack of considering multiple factors. Accordingly, our research project aims to address the following research goal (RG):

RG: *Develop and evaluate an algorithm that considers multiple DR-objectives.*

With this algorithm several stakeholders can benefit such as users as they can reach several goals (e.g., cost minimization and welfare maximization), energy providers as they can give more incentives to users and achieve a more efficient generation or predictability, energy grids can get more stable (e.g., ensuring reliability and market performance) and flexible on contingencies, appliances (e.g., EVs) and infrastructure (e.g., storages and PVs) will be embedded in a more fertile way.

We already did first steps in designing such an algorithm in a former study [13]. However, further research is needed, to implement and evaluate it. Therefore, our paper is structured as follows: After describing the optimization problem (Sect. 2) the research process, adopted from the “algorithm engineering” (AE) approach, is described (Sect. 3). Afterwards, the algorithm of our former study is described [13] (Sect. 4). According to the AE approach, we implement the algorithm (Sect. 5), which means the instantiation of individual parameters within the algorithm. Afterwards, we derive an experimental setting and simulated four test cases with input data from the DR field (Sect. 6). We finish with a conclusion, discuss limitations and give an outlook for further research (Sect. 7).

2 Multi-objective Optimization in the DR Field

2.1 Related Work

To identify related research, we conducted a literature search using Google Scholar and AISeL (covering IEEE and ScienceDirect), for a wide range of potentially relevant articles. Additionally, we analyzed research mentioned in literature reviews of [14–16]. Here, especially the supplier side has been in the focus (e.g., Trading-Agent-Competition [17]). On the consumer side, to the best of our knowledge, no appropriate

solution could be found. Addressed goals are heterogeneous, such as cost reduction (e.g., [18–20]) or peak load reduction (e.g., [21–23]). Some algorithms try to achieve two goals simultaneously (e.g., reducing costs while respecting users’ comfort, e.g., [24–27]). These goals are often realized based on an externally given cost function of an energy supplier and the comfort level, that is considered as an additional condition to the objective function. [28] consider multiple goals, however, these are not suitable for our research goal, as multiple and difficult to solve optimization functions are used. Algorithms in particular try to combine goals, for example, to minimize costs and maximize the comfort level of users (e.g., [12, 25, 27]). Batchu and Pindoriya [29] identified four indicators, addressing major DR objectives: lower energy consumption, peak load reduction, load profile flattening and cost reductions. The optimization problem is constructed with multiple objective functions, optimizing towards load profile flattening, cost reduction, comfort maximization and peak load reduction. Multiple objective functions mostly return in a complex problem and solving is difficult. The Berkeley Lab ([28]) addressed multiple goals, however, the resulting optimization problem will get very complex and challenging to solve. They therefore propose a Greedy algorithm to find a solution.

2.2 DR Optimization Problem

The main goal of our study is to consider multiple objectives in DR algorithms. A constrained multi-objective optimization problem (CMOP) (already used in the DR field [29]) is defined as follows (e.g., [12, 30]) (x = vector of inputs):

$$F(x) = [f_1(x), f_2(x), \dots, f_k(x)] \tag{1}$$

Under the additional conditions ($g()$ and $h()$ = conditional functions):

$$g_i(x) = 0 \quad i = 1, \dots, m \tag{2}$$

$$h_j(x) < 0 \quad j = 1, \dots, n \tag{3}$$

In order to derive a single optimization function model, several constraints regarding the appliances [31] have to be considered, to meet the application area. We therefore choose the optimization model of [32], based on [31], which aims at reducing the costs of the overall system with the help of a ToUP, which can be exchanged to several other DP alternatives. Our model uses the following variables: Let N be all considered living units and A_n be all the appliances of living unit $n \in N$ and ω be the sample rate of the discrete model (number of time periods) over one day. Moreover, let $x^h = \sum_{n \in N} \sum_{a \in A_n} h_{n,a}^h$ with $h \in Z = \{0, 1, \dots, \omega\}$ be the sum of all appliances $a \in A_n$ of all living unit $n \in N$ in the timeslot h .

Let $l_{n,a}^k$ be the load profile in a local time interval $k \in T_1 = \{0, 1, \dots, \delta_{n,a}\}$, $l_{n,a} = \sum_{k \in T_1} l_{n,a}^k$ the load sum and $\delta_{n,a}$ the length of load $a \in A_n$. In doing so, we can transform a given horizontal and inseparable load profile from its local time interval T_1 to the global

one T (7) through shifting the whole T_l by an appropriate constant $m_{n,a}$, i.e., $h_k = k + m_{n,a}$ with $0 \leq m_{n,a} \leq \omega - \delta_{n,a}$.

Furthermore, let $\gamma_{n,a}^{h,\min}$ be the min and $\gamma_{n,a}^{h,\max}$ be the max borders for a load $x_{n,a}^h$ with $h \in T, a \in A_n$ so we can specify, in which borders the intensity of load a can be shifted, i.e., $\gamma_{n,a}^{h,\min} \leq x_{n,a}^h \leq \gamma_{n,a}^{h,\max}$ (8). We note that the given load profiles have to satisfy the inequality $\gamma_{n,a}^{h,\min} \leq I_{n,a}^{h_k} \leq \gamma_{n,a}^{h,\max}$ for all $k \in T_l$ to get a feasible solution.

Let $\alpha_{n,a}$ be the starting and $\beta_{n,a}$ be the ending time slot for an appliance a . with the restricted time interval $T [\alpha_{n,a}, \beta_{n,a}]$ (9). Note that the interval length must be at least the length of the load profile $\delta_{n,a}$. to get a feasible solution, i.e., $\beta_{n,a} - \alpha_{n,a} \geq \delta_{n,a}$.

In order to turn on and off the constraint i for each appliance a individually (e.g., an EV has other constraints than a washing machine) let $c_a^i \in \{0, 1\}$ be a binary variable that shows if a constraint is turned on ($c_a^i = 1$) or not ($c_a^i = 0$) for the appliance a .

The objective function describes the total cost (4), while the cost in a time slot h is a function depending on h and the total load x^h i.e., $c^h = c(h, x^h) * x^h$.

$$\min_{x_{n,a}^h} c = \sum_{h=0}^{\omega} c^h = \sum_{h=0}^{\omega} c(h, x^h) * x^h \tag{4}$$

$$\sum_{h=0}^{\omega} x_{n,a}^h = I_{n,a} \quad \forall n \in N, a \in A_n \tag{5}$$

$$(x_{n,a}^{h_k} - I_{n,a}^k) * c_a^1 = 0 \quad \forall k = 0, \dots, \delta_{n,a}, \forall n \in N, a \in A_n \tag{6}$$

$$\left. \begin{aligned} (x_{n,a}^h - \gamma_{n,a}^{h,\max}) * c_a^2 &\leq 0 \\ (\gamma_{n,a}^{h,\min} - x_{n,a}^h) * c_a^2 &\leq 0 \end{aligned} \right\} \forall n \in N, a \in A_n \tag{7}$$

$$\left(\sum_{h=\alpha_{n,a}}^{\beta_{n,a}} x_{n,a}^h - I_{n,a} \right) * c_a^3 = 0 \quad \forall n \in N, a \in A_n \tag{8}$$

$$x_{n,a}^h \geq 0, c_a^i \in \{0, 1\} \quad \forall h \in T, n \in N, a \in A_n, i = 1, 2, 3 \tag{9}$$

3 Research Methodology

Our research project follows the Design-Science-Research (DSR) paradigm [33], aiming to develop a new DR algorithm. To design an algorithm, we basically found two different approaches: (a) “algorithm theory” (AT) and (b) “algorithm engineering” (AE) [34]. AT focusses on the formulation of the mathematical model and afterwards proofs that it finds an (optimal) solution in a guaranteed time. AE addresses an application area, which gets more important and complex nowadays (e.g., data volume increases and real time computing). AE does not aim to proof a guaranteed time and optimization for all possible inputs but rather to guarantee a good performance and result for common inputs of the application area. Because the DR context (application area), is a complex field (e.g., stated as a *wicked problem* [35]) and the evaluation in

DSR should focus the real world problem, respectively the application area (e.g., [36]), we assume the AE approach as more suitable (cf., [37]).

The AE approach consists of five different (main)steps (see Fig. 1): *modelling* the application area as basis for the *design* phase, followed by the *analysis* of the designed algorithm, *implementation* of the algorithm and finally an *experiment*. All of these steps have a strong relationship to the application, receiving application depending information (e.g., input data or the mathematical model) or contributing to the application (e.g., delivering a library to use or giving a performance guarantee).

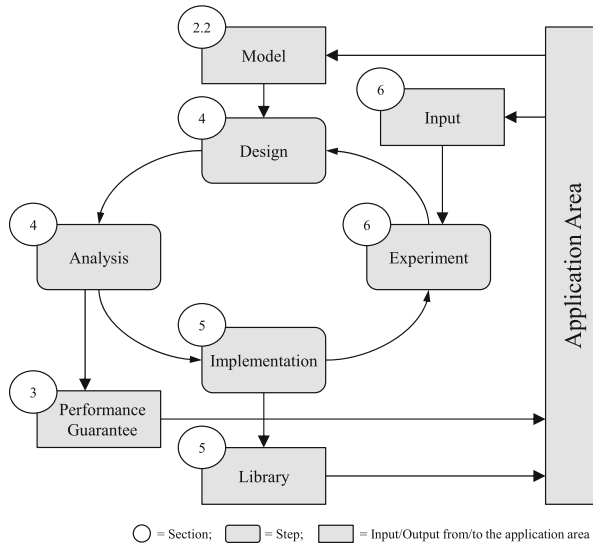


Fig. 1. Research design (cf. [34]).

4 Design and Analysis of the Algorithm

The optimization model in Sect. 2.1 only aims at cost reduction. Hence, we need to identify suitable DR objectives to design an optimization model. [29] identified several indicators that need to be integrated in the optimization model, as these indicators match with most of the DR objectives. These additional objectives need to be included in the optimization function (e.g., done by [29] in a generalized form for a single home). The authors use multiple optimization functions, which results in an even more complex problem and the algorithm gets inefficient, for example, regarding the calculation time [32]. We therefore need less objective functions.

One approach within the stated model could be ‘penalty costs’. The basic idea is that all objectives influence a cost factor, which then, based on a kWh price, gives the costs in a certain timestamp for placing the load. The objective function is formulated

in the following way, over the optimization horizon: $\min_{x_{n,a}^h} c = \sum_{h=0}^{\omega} c(\theta) * x^h$, where θ is the vector of additionally needed information to calculate the penalty costs [13].

First of all, the timeslot (time-of-use) is needed to consider time dependent costs (e.g., ToUP). In order to achieve a *flattening* of the load profile, the fluctuation needs to be minimized. Therefore, we implement rising penalty costs for deviations from the arithmetic mean of the last X timestamps. This means, if the placement of a certain load in timeslot Y rises or lowers the consumption in that timeslot more than a certain percentage compared to the arithmetic mean of the last X timeslots, the cost factor is raised and penalty costs are added. The same procedure is chosen to integrate a more effective *peak load reduction*. If the placement in timeslot Y will create a new peak, the cost factor should be raised by a certain percentage.

At this point, two upcoming questions must be answered: how can the additional information be gathered (see following section), and how do the penalty costs look like. The penalty costs need certain parameters, which need to be instantiated (filled with concrete values). For example, how much are the penalty costs increasing when creating a new peak? However, because the selection of the concrete values does not influence the performance of the algorithm according its runtime, we do this in the following analysis phase. The objective function now looks the following:

$$\min_{x_{n,a}^h} c = \sum_{h=0}^{\omega} c^h = \sum_{h=0}^{\omega} c(\theta) * x^h \tag{10}$$

Both the function $c(h)$ and the vector θ still need to be defined. θ needs to involve information about the timestamp h , the load in this timestamp x^h , the arithmetic mean of the last X timestamps $MEAN(h - X)$ and the actual peak load $PEAK$. Moreover, parameters for (a) the cost function (a ToUP with two different intervals is chosen (low price (LP) and high price (HP))) with additional information, such as costs in different intervals and interval lengths, (b) increasing cost rate ρ for $\sigma\%$ of deviation from the mean, and (c) increasing cost rate τ for a new peak are needed.

$$\begin{aligned} c(\theta) &= c(h, x^h, MEAN(h - 8), PEAK) \\ &= c_1 \rightarrow \begin{cases} c = LP : 0 \leq h \leq a; b \leq h \leq \omega \\ c = HP : a < h < b \end{cases} \\ &+ c_2 \rightarrow \begin{cases} c = c * \rho : 0 * \sigma \leq < 1 * \sigma \\ c = c * \rho : 1 * \sigma \leq \frac{|x^h - MEAN(h - X)|}{x^h} < 2 * \sigma \\ c = c * \rho : 2 * \sigma \leq < 3 * \sigma \\ \dots & \dots \end{cases} \\ &+ c_3 \rightarrow \begin{cases} c = c * \tau : x^h = PEAK \\ c = c : \textit{else} \end{cases} \end{aligned} \tag{11}$$

We now need to investigate, if the algorithm terminates and how the runtime is. Both depend on the chosen solution strategy. However, if a suitable heuristic (e.g., Greedy heuristic) is chosen, we can guarantee termination. Calculating the runtime, n

loads need to be sorted, resulting in an average runtime of $O(n \cdot \log(n))$ and $O(n^2)$ in the worst case (e.g., quicksort). Moreover, our placement algorithm needs to look at m timeslots and calculate the costs. The runtime therefore is $O(n \cdot m)$. The overall runtime is then $O(m + n \cdot \log(n))$ for the average case and $O(n \cdot m + n^2)$ for the worst case. As we can expect the average case most likely or we can obligate the users to deliver the loads pre-sorted, we can expect an average case. As we are in the AE setting (application area is more important than the worst case), the guaranteed performance is $O(n \cdot m + n \cdot \log(n)) = O(n \cdot (m + \log(n)))$.

5 Implementation

5.1 Instantiation of the Parameters

Several parameters need to be instantiated in order to implement the algorithm. These parameters are divided into: (a) factors increasing the costs and (b) boundaries in which the cost factor is increased. (a) involves the parameters LP, HP, ρ and τ , and (b) the parameters a , b , X and σ . Because different instantiations are possible, we need to investigate different combinations of parameter instantiations. As a first step, we highlight aspects to be regarded during the instantiation. Afterwards, we will simulate different combinations to evaluate which instantiation will fit best (see Sect. 6).

- (a) *Low- and High-Price (LP, HP)*: The price of one kWh is ~ 0.3 €. In a high price period it increases (+ ~ 3 ct) and in a low pricing period it falls ($- \sim 4$ ct).
- (a) *Cost-increase per σ -load-profile-flattening (ρ)*: If the load profile is flattening too much, the cost factor is increased by ρ per σ % of load profile flattening, meaning variation from the arithmetic mean of the last X timestamps.
- (a) *Cost-increase if a new peak is created (τ)*: If a new peak is created, we need to determine, how much the costs should be increased. As these peaks are very critical, the factor should be higher than for prohibiting load fluctuation.
- (b) *Pricing intervals (a , b)*: These intervals define the high and the low pricing times. Usually these are set in the morning and in the evening (low price, e.g. 0–5 h and 22–24 h) as well as during the day for the high price (e.g., 5–22 h).
- (b) *Timestamps for the arithmetic MEAN (X)*: Initial simulations indicated 8 timestamps (two hours) suitable for rating the load profile flattening.
- (b) *Percentages of load profile flattening for cost increase (σ)*: Percentage σ of deviation from the mean. Each σ percent the costs are increased.

5.2 Needs for Additional Information

The initial model requires a cost function, depending on the timeslot and the amount of consumption in this timeslot. This information is suitable for most pricing schemes such as ToUP and RTP. With the new objective function, further information about the arithmetic mean of the last 8 timeslots and the highest peak is needed. Depending on which control scheme (e.g., [7]), this information can be given on different ways:

- *Autonomous control*: This control scheme has no communication and a local decision making. The local controller can only access local information. The optimization is done based on this information, that should be accessible (no security and privacy concerns have to be dealt with) [38].
- *Direct control*: In this case (one or two-way communication) the centralized controller has all information needed and the decision making about the load shifting is made directly as well. Privacy concerns need to be regarded.
- *Indirect control and transactional control*: The decision making is done locally. However, because of the one or two-way communication infrastructure, the information needed can be send from the central controller or the other users. Besides the privacy concerns, it has to be argued, which information is used: only the local information of each user or the overall information from the grid.

5.3 Possible Solving Strategies

To solve the stated CMOP, a suitable method is needed, for example evolutionary algorithms, which can find good approximations for CMOPs [12]. Other possibilities are for example Greedy, Hill-Climbing or Branch-and-Bound algorithms (e.g., [39, 40]). The Greedy heuristic is a method that is relative easy to implement with a guaranteed runtime and known in the DR field. As especially the runtime needs attention here, the Greedy heuristic is chosen. Other heuristics are possible as well.

Solving the optimization problem with the Greedy heuristic can be carried out in two different ways. The “classic” Greedy algorithm—in the DR field—sorts all shiftable loads either downward or upward according their (summarized) consumption. This means, the biggest or lowest load is picked first. The load is placed in the “best” timeslot (see Table 1). Regardless which sorting strategy is selected, we need to identify the best timeslot. In former Greedy applications, this is, for example, the place with the lowest energy consumption or the timeslot with the highest energy generation. In our algorithm, the best placement is threefold: (a) cheapest place according the cost-function, (b) place with the lowest resulting fluctuation of the load profile, and (c) place where no new peak is created. An advantage of our approach is that the functionality remains more or less the same compared to the “classic” Greedy. Only the identification of the best placement changes.

Table 1. (Classic) greedy pseudocode.

Start	Take maximum/minimum load	
02		Search the cheapest place for the load
03		Calculate consumption – generation
04		Find lowest result
05		Place the load at this place
End	Take the next load and start again	

6 Experiment

6.1 Input Data

Conducting the experiment according the stated AE approach needs suitable input data from the application area. Recorded data from naturalistic living units (analyzed for example in [41, 42]) or artificial data, which has a sufficient quality ([43, 44]) seems both possible. In the naturalistic data case, we must add additional information, for example, the usage time of appliances. For generating artificial data, the LoadProfileGenerator (LPG) [45] can be used that has predefined load profiles and uses a behavior model for each user in a living unit to simulate the data.

As artificial data has a good quality compared to recorded data [43], we use the LPG. We choose a household with a family (two children, man at work) and a washing machine, a dryer and a dishwasher as deferrable appliances. The household has a ToUP with 21 cent per kWh during off-peak hours (0–5 and 22–24 h) and 39 cents in peak hours (5 to 22 h). As instantiations for our algorithm, we decided to choose a couple of variations, which can be seen in the Table 2. To enable a good comparison of our algorithm, we implemented further algorithms, aiming at different objectives: the first only aims at cost minimization that is achieved by the ToUP (here called ToU). The second just aims to realize a load profile flattening. It uses a notional convex cost function (as used for example in [46], here called ECF).

Table 2. Instantiation of Parameters.

Variation	LP	HP	a	b	ρ	σ	τ
New1	21 Cent	39 Cent	5 h	22 h	10%	0.1	2
New2					20%	0.1	
New3					10%	0.2	
New4					20%	0.2	
New5					5%	0.1	
New6					5%	0.2	

6.2 Experiment Results

The simulation shows, that the runtime was as expected. For analyzing the different algorithms, we use different indicators, which can be seen in Table 3. The results in Table 4 indicate that ECF and ToU either realize a flattened load profile or a good cost minimization. Depending on the instantiation of our algorithm, we achieve a cost minimization slightly under the TuO but a more flattened load profile (except instantiation 2) or even a more cost benefits while the load profile is more fluctuation then the TuO one. Same results can be seen when looking at the peak load reduction.

Overall it can be seen that the second instantiation (New2) performs best according the costs, however the load profile is more flattening and a new peak is created. Instantiation New4 achieves nearly the same savings compared to the ToU and the load profile is nearly as flattened as in the ECF case. Hence, New4 performs best, as it achieves a result near the best performing algorithms according costs, PAR and MSE.

Table 3. Identified indicators.

Indicators [unit]	Description
Mean (e.g., [47]) [kWh]	Average consumption (in kWh) over the day
Peak-to-Average-Ratio (PAR) (e.g., [48]) [kWh]	Highest peak from the average (positive or negative) over the day
Mean-Squared-Error (MSE) (e.g., [23]) [kWh ²]	The arithmetic middle of the squared error during the day. Often also indicted as Root-Mean-Squared-Error (RMSE)
Costs (e.g., [49]) [€] or [\$]	How much must be paid? Pricing scheme is needed
Savings (e.g., [50]) [€/\$] or [%]	How much savings can be realized comparing to no DR?

Table 4. Experiment results.

Factor	No DR	ECF	ToU	New1	New2	New3	New4	New5	New6
Mean	0,2328	0,2328	0,2328	0,2328	0,2328	0,2328	0,2328	0,2328	0,2328
PAR	0,9967	0,7560	1,1393	0,8355	1,0646	0,8355	0,7791	0,8355	0,8355
MSE	6,1699	3,7535	4,6109	4,0675	4,9797	4,0675	3,7891	4,0675	4,0675
Costs	8,0143	7,5550	7,0358	7,0475	6,7071	7,0475	7,0475	7,0475	7,0475
Savings	0	0,0573	0,1221	0,1206	0,1631	0,1206	0,1206	0,1206	0,1206

7 Contribution, Limitations, and Further Research

In this study, an evaluated, multi-objective DR algorithm for the residential context is proposed. In doing so, several stakeholders can benefit such as users as they can reach several goals (e.g., cost minimization), energy providers as they can give more incentives and achieve a more efficient generation and predictability, energy grids can increase reliability and react more effectively and flexible on contingencies.

Although we derived helpful insights for practice and research, our work is not free of limitations. First, we cannot ensure that we found all relevant factors. However, even if a new factor has to be taken into account, the optimization problem and the cost function can be extended. Second, we cannot assure that the formulated optimization problem or the solving strategy (Greedy heuristic) are the best or optimal solution. As we did not state to formulate or find the best performing solution, but a first iteration towards this, we assume the research question answered (to this point).

For further research, we plan to evaluate the algorithm in the field (ex post, naturalistic evaluation). Together with an industry partner who develops a Home Energy Management System, we want to install and test this algorithm in a real-world setting. Moreover, we plan to investigate, if other instantiations and on which conditions these instantiations perform best. This field, also regarded as “hyper-parameter-optimization” needs further investigation (e.g., [51]).

References

1. Lawrence, T.M., Watson, R.T., Boudreau, M.-C., Mohammadpour, J.: Data flow requirements for integrating smart buildings and a smart grid through model predictive control. *Procedia Eng.* **180**, 1402–1412 (2017)
2. Seidel, S., Recker, J., vom Brocke, J.: Sensemaking and sustainable practicing: functional affordances of information systems in green transformations. *Manag. Inf. Syst. Q.* **37**, 1275–1299 (2013)
3. Siano, P.: Demand response and smart grids—a survey. *Renew. Sustain. Energy Rev.* **30**, 461–478 (2014)
4. Hu, Q., Li, F.: Hardware design of smart home energy management system with dynamic price response. *IEEE Trans. Smart Grid.* **4**, 1878–1887 (2013)
5. Jovanovic, R., Bousseham, A., Bayram, I.S.: Residential demand response scheduling with consideration of consumer preferences. *Appl. Sci.* **6**, (2016)
6. Koolen, D., Sadat-Razavi, N., Ketter, W.: Machine learning for identifying demand patterns of home energy management systems with dynamic electricity pricing. *Appl. Sci.* **7**, 1160 (2017)
7. Kosek, A.M., Costanzo, G.T., Bindner, H.W., Gehrke, O.: An overview of demand side management control schemes for buildings in smart grids. In: 2013 IEEE International Conference on Smart Energy Grid Engineering. SEGE (2013)
8. Merkert, L., Harjunkoski, I., Isaksson, A., Säynevirta, S., Saarela, A., Sand, G.: Scheduling and energy – industrial challenges and opportunities. *Comput. Chem. Eng.* **72**, 183–198 (2015)
9. Steen, D., Le, T., Bertling, L.: Price-based demand-side management for reducing peak demand in electrical distribution systems – with examples from gothenburg. Chalmers Publication Library (CPL) (2012)
10. Hillemacher, L.: Lastmanagement mittels dynamischer Strompreissignale bei Haushaltskunden. (2014)
11. Vidal, A.R.S., Jacobs, L.A.A., Batista, L.S.: An evolutionary approach for the demand side management optimization in smart grid. In: 2014 IEEE Symposium on Computational Intelligence Applications in Smart Grid (CIASG), pp. 1–7 (2014)
12. Salinas, S., Li, M., Li, P.: Multi-objective optimal energy consumption scheduling in smart grids. *IEEE Trans. Smart Grid.* **4**, 341–348 (2013)
13. Behrens, D., Schoormann, T., Knackstedt, R.: Developing an algorithm to consider multiple demand response objectives. *Eng. Technol. Appl. Sci. Res.* (2018)
14. Al-Sumaiti, A.S., Ahmed, M.H., Salama, M.M.A.: Smart home activities: a literature review. *Electr. Power Compon. Syst.* **42**, 294–305 (2014)
15. Balijepalli, V.S.K.M., Pradhan, V., Khaparde, S.A., Shereef, R.M.: Review of demand response under smart grid paradigm. In: ISGT2011-India, pp. 236–243 (2011)
16. Gerwig, C., Behrens, D., Lessing, H., Knackstedt, R.: Demand side management in residential contexts - a literature review. In: Lecture Notes in Informatics, pp. 93–107 (2015)
17. Ketter, W., Collins, J., Reddy, P.P., Flath, C.M.: The Power Trading Agent Competition. Social Science Research Network, Rochester (2011)
18. Ali, S.Q., Maqbool, S.D., Ahamed, T.P.I., Malik, N.H.: Pursuit algorithm for optimized load scheduling. In: 2012 IEEE International Power Engineering and Optimization Conference Melaka, Malaysia, pp. 193–198 (2012)
19. Keerthisinghe, C., Verbič, G., Chapman, A.C.: Evaluation of a multi-stage stochastic optimisation framework for energy management of residential PV-storage systems. In: 2014 Australasian Universities Power Engineering Conference (AUPEC), pp. 1–6 (2014)

20. Zhao, W., Cooper, P., Perez, P., Ding, L.: Cost-driven residential energy management for adaptation of smart grid and local power generation (2014)
21. Huang, Y., Mao, S., Nelms, R.M.: Smooth electric power scheduling in power distribution networks. In: 2012 IEEE Globecom Workshops, pp. 1469–1473 (2012)
22. McNamara, P., McLoone, S.: Hierarchical demand response for peak minimization using dantzig–wolfe decomposition. *IEEE Trans. Smart Grid.* **6**, 2807–2815 (2015)
23. Verschae, R., Kawashima, H., Kato, T., Matsuyama, T.: A distributed coordination framework for on-line scheduling and power demand balancing of households communities. In: 2014 European Control Conference (ECC), pp. 1655–1662 (2014)
24. De Angelis, F., Boaro, M., Fuselli, D., Squartini, S., Piazza, F., Wei, Q., Wang, D.: Optimal task and energy scheduling in dynamic residential scenarios. In: Wang, J., Yen, G.G., Polycarpou, M.M. (eds.) *ISNN 2012. LNCS*, vol. 7367, pp. 650–658. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-31346-2_73
25. Bassamzadeh, N., Ghanem, R., Kazemitabar, S.J.: Robust scheduling of smart appliances with uncertain electricity prices in a heterogeneous population. *Energy Build.* **84** (2014)
26. Kim, S.-J., Giannakis, G.B.: Efficient and scalable demand response for the smart power grid. In: 2011 4th IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP), pp. 109–112 (2011)
27. Song, L., Xiao, Y., van der Schaar, M.: Non-stationary demand side management method for smart grids. In: 2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 7759–7763 (2014)
28. Berkeley Lab: Distributed Energy Resources Customer Adoption Model (DER-CAM) | Building Microgrid. <https://building-microgrid.lbl.gov/projects/der-cam>
29. Batchu, R., Pindoriya, Naran M.: Residential demand response algorithms: state-of-the-art, key issues and challenges. In: Pillai, P., Hu, Y.F., Otung, I., Giambene, G. (eds.) *WiSATS 2015. LNCS*, vol. 154, pp. 18–32. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-25479-1_2
30. Cortés-Arcos, T., Bernal-Agustín, J.L., Dufo-López, R., Lujano-Rojas, J.M., Contreras, J.: Multi-objective demand response to real-time prices (RTP) using a task scheduling methodology. *Energy* **138**, 19–31 (2017)
31. Behrens, D., Schoormann, T., Knackstedt, R.: Towards a taxonomy of constraints in demand-side-management-methods for a residential context. In: Abramowicz, W. (ed.) *BIS 2017. LNCS*, vol. 288, pp. 283–295. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-59336-4_20
32. Behrens, D., Ruether, C., Schoormann, T., Ambrosi, K., Knackstedt, R.: Effects of constraints in residential demand-side-management algorithms - a simulation-based study. In: *International Conference on Operation Research*. Springer, Berlin (2017)
33. Hevner, A.R., March, S.T., Park, J., Ram, S.: Design science in information systems research. *Manag. Inf. Syst. Q.* **28**, 75–105 (2004)
34. Sanders, P.: Algorithm engineering – an attempt at a definition. In: Albers, S., Alt, H., Näher, S. (eds.) *Efficient Algorithms. LNCS*, vol. 5760, pp. 321–340. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-03456-5_22
35. Ketter, W., Peters, M., Collins, J., Gupta, A.: competitive benchmarking: an is research approach to address wicked problems with big data and analytics. *Manag. Inf. Syst. Q.* **40**, 1057–1080 (2016)
36. Füller, K., Ramanath, R., Böhm, M., Krcmar, H.: Decision support for the selection of appropriate customer integration methods. *Wirtsch. Proc.* **2015** (2015)
37. Nieffe, A., Tröschel, M., Sonnenschein, M.: Designing dependable and sustainable smart grids – how to apply algorithm engineering to distributed control in power systems. *Environ. Model Softw.* **56**, 37–51 (2014)

38. Mohsenian-Rad, A.-H., Wong, V.W., Jatskevich, J., Schober, R., Leon-Garcia, A.: Autonomous demand-side management based on game-theoretic energy consumption scheduling for the future smart grid. *IEEE Trans. Smart Grid*, **1**, 320–331 (2010)
39. Nelder, J.A., Mead, R.: A simplex method for function minimization. *Comput. J.* **7**, 308–313 (1965)
40. Land, A.H., Doig, A.G.: An automatic method of solving discrete programming problems. *Econometrica* **28**, 497–520 (1960)
41. Monacchi, A., Egarter, D., Elmenreich, W., D’Alessandro, S., Tonello, A.M.: GREEND: an energy consumption dataset of households in Italy and Austria. *ArXiv14053100* (2014)
42. Behrens, D., Schoormann, T., Knackstedt, R.: Datensets für Demand-Side-Management – Literatur-Review-Basierte Analyse und Forschungsagenda. In: Mayr, H.C., Pinzger, M. (eds.) *Lecture Notes in Informatics* (2016)
43. Cao, H.Á., Beckel, C., Staake, T.: Are domestic load profiles stable over time? An attempt to identify target households for demand side management campaigns. In: *IECON 2013 - 39th Annual Conference of the IEEE Industrial Electronics Society*, pp. 4733–4738 (2013)
44. Hoogsteen, G., Molderink, A., Hurink, J.L., Smit, G.J.M.: Generation of flexible domestic load profiles to evaluate demand side management approaches. In: *2016 IEEE International Energy Conference (ENERGYCON)*, pp. 1–6 (2016)
45. Noah Pflugradt: LoadProfileGenerator. <http://www.loadprofilegenerator.de/>
46. Behrens, D., Gerwig, C., Knackstedt, R., Lessing, H.: Selbstregulierende Verbraucher im Smart Grid: Design einer Infrastruktur mit Hilfe eines Multi-Agenten-Systems. In: *Proceedings of the Multikonferenz Wirtschaftsinformatik 2014* (2014)
47. Miao, H., Huang, X., Chen, G.: A genetic evolutionary task scheduling method for energy efficiency in smart homes. *Int. Rev. Electr. Eng.* **7**, 5897–5904 (2012)
48. Soliman, H.M., Leon-Garcia, A.: Game-theoretic demand-side management with storage devices for the future smart grid. *IEEE Trans. Smart Grid*, **5**, 1475–1485 (2014)
49. Alam, M.R., St-Hilaire, M., Kunz, T.: Cost optimization via rescheduling in smart grids—a linear programming approach. In: *2013 IEEE International Conference on Smart Energy Grid Engineering (SEGE)*, pp. 1–6 (2013)
50. Maqbool, S.D., Ahamed, T.P.I., Ali, S.Q., Pazheri, F.R., Malik, N.H.: Comparison of pursuit and ϵ -greedy algorithm for load scheduling under real time pricing. In: *2012 IEEE International Conference on Power and Energy (PECon)*, pp. 515–519 (2012)
51. Bergstra, J.S., Bardenet, R., Bengio, Y., Kégl, B.: Algorithms for hyper-parameter optimization. In: *Advances in Neural Information Processing Systems* (2011)



Toward Resilient Mobile Integration Processes

Daniel Ritter^(✉) and Manuel Holzleitner

SAP SE, Dietmar-Hopp-Allee 16, 69190 Walldorf, Germany
{daniel.ritter,manuel.holzleitner}@sap.com

Abstract. The widespread use of smart mobile devices and the digital transformation foster context-aware applications (short: apps) that bridge the gap between the socio-technical and business worlds in our everyday lives. While the number of proprietary, mobile app solutions is increasing, a mobile integration system architecture, e. g., connecting mobile and cloud apps, and studies of the trade-off between resource-limits and resilient service qualities on mobile computing platforms are missing.

In this work we define the core components of a resilient, resource-aware, mobile enterprise application integration (EAI) system, which requires extensions beyond the current non-mobile architectures. We evaluate the proposed architecture and especially its runtime resource and monitoring characteristics in form of a prototypical mobile realization.

Keywords: Context-aware applications
Enterprise application integration · Mobile integration system
Resilience patterns

1 Introduction

Current smart mobile devices are in the center of our everyday lives by providing a huge set of applications (called apps). In a process-like manner, these apps connect the user's context (e. g., location, car engine status) with cloud services (e. g., ERP predictive maintenance), which are called context-aware applications (e. g., [20, 22]). To build more advanced, context-aware mobile applications, there has been some recent work, e. g., on mobile process [20] and service engines [19] for data collection scenarios. However, the current mobile apps have two shortcomings: (1) the apps are isolated and communicate via proprietary RPC implementations, leading to n-square connectivity and data format problems [11] that are solved by enterprise application integration (EAI) [11], and (2) the apps might neglect the trade-off between communication qualities (e. g., low latency, no data loss) and resource-awareness or stability (e. g., CPU, battery consumption) as also identified as challenge in [1]. Figure 1(a) shows the predominant integration situation in current mobile app solutions (cf. (1)). Most of the apps are directly connected to multi-cloud apps or services [24]. The redundant communication layers of these apps (cf. marked as *com.*) are proprietary, and thus

have varying integration qualities (e. g., security) and increase the overall memory footprint on the device (e. g., multiple RPC implementations). The inter-app communication via IPC on the same device is rarely used and the underlying mobile OS concepts differ largely by their vendors. The mobile platforms offer services (e. g., storage, network) and context information (e. g., contacts, GPS). As a solution, we argue that a central integration system on the mobile devices, shown in Fig. 1(b), solves the problems from (1), reduces the communication logic within the applications and (2) allows for a central regulation of the resource consumption for all communication tasks. In contrast to existing service (e. g., [20]) and automation solutions (e. g., IFTTT¹) EAI systems allows for complex routing and transforming patterns (e. g., [16,17]), whose compositions we call integration scenarios.

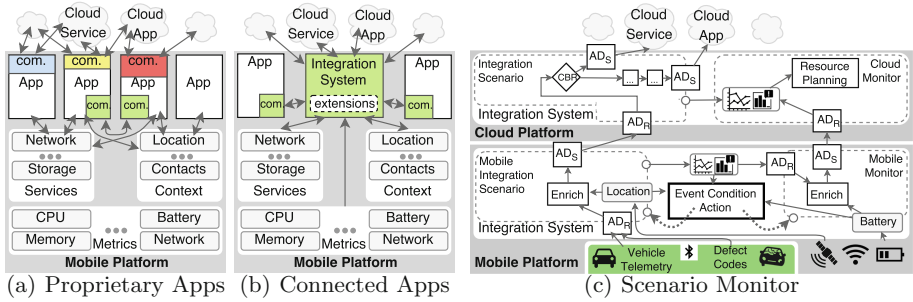


Fig. 1. Current vs. desired EAI solution and monitor on mobile computing systems

For example, Fig. 1(c) denotes a vehicle logistics integration scenario, which requires central integration logic as part of a cloud solution and a local, context-aware scenario part on the device. Essentially, the local part collects data from the environment (e. g., vehicle speed, and defect codes) in a high frequency, and sends aggregates to the cloud solution with a configurable timing (e. g., every 30 s). The cloud solution adds master data to the incoming messages, determines the receiving analytical applications and transforms the content accordingly. Furthermore, monitoring data from the EAI system (e. g., number of processed messages or threads) and the mobile platform (e. g., resource capacities and usage) are transferred and displayed in specific integration monitors according to [16] (e. g., message, resource monitors). However, the current EAI architectures are not designed for mobile deployments and miss resource-aware features, leading to the research questions, addressed in this work:

- Q1 Which device data are relevant for a stable integration solution and what are suitable reactions on exceptional situations (e. g., offline, low battery)?
- Q2 Which mechanism allows for building a context-aware integration system with a resource-aware monitoring for multiple devices and applications?

¹ IF This Than That, visit. 02/2018: <https://ifttt.com/>.

- Q3 Which architecture components are currently missing and what impact do they have on the overall integration scenario runtime?
- Q4 Which side-effects do the monitoring and resilience aspects have on the message processing and the device itself?

To answer these questions, we determine and classify mobile device data and countermeasures (cf. Q1) tailored to mobile devices in Sect. 2. From the resulting impact analysis, we derive and define an adaptable mechanism in Sect. 3 focusing on resilient monitoring (cf. Q2), however, applicable to the runtime system. The term *resilient* was also used in [1] in the context of component failures, while we extend it for mobile devices as a characteristic to withstand several exceptional context or resource-events (cf. Q1), to continue to operate normally as long as possible. Then we describe the required extensions to the existing EAI architectures in Sect. 4 (cf. Q3) and evaluate our approach in Sect. 5 (cf. Q4). We discuss related work in Sect. 6 and conclude in Sect. 7.

2 Mobile Integration System Monitoring

In this section, we classify the common device data of current smart mobile devices, denoting resource or context information and potential issues, and set them into context to the countermeasures treating these issues in form of exception and integration adapter patterns from existing integration pattern catalogs [16, 17].

2.1 Mobile Device Data and Impact Classification

We differentiate between the device data (e. g., from device specifications) and the data of EAI systems (e. g., #msgs/sec) [11]. Due to our focus on integration system monitoring, we assume that the mobile monitor gathers configurable metrics of the local integration system for local (i. e., on device) or remote visualization and alerting (e. g., cloud monitor [16]). Hence, we concentrate on the context and resource data, which have impact on the stability of the system (cf. Q1) and assume that all app metrics can be subsumed by resource data (e. g., increasing #msgs/sec \rightarrow CPU usage or temperature). We distinguish between the capacity and the actual usage. Similarly, the EAI system data could be classified.

Context Data and Device Resources. Although our approach is not limited to the *network* (cf. Fig. 1(b)), it is the only context data that we discuss in this work. The network on a mobile device requires wireless network access either in form of bluetooth, a WIFI module or mobile data via a GSM module. This also determines the three states that a device can be in: {offline, online, mobile data} for the metric that we call *mobile data state* as metric (M0). Note that there is other data like SSID (WIFI provider) or RSSI (for signal strength), which we do not consider in this work, however, could be incorporated, if required.

With respect to resources, a mobile device is currently constituted according to a van-Neumann architecture with a CPU, random access memory (RAM), and a (solid state) hard drive. When combining these components with their general characteristics {usage, temperature} we come to the following metrics: *CPU usage* (M1), *CPU temp* (M2), *Free Memory* (M3), *Hard drive usage* (M4). In addition, mobile devices have a battery as power supply, leading to the metrics: *Battery level* (M5), *Battery temperature* (M6).

Metrics Impact Analysis. To determine the impact of each metric we categorize them according to the severity from [10]: critical, urgent, intervention, recoverable. Thereby we consider CPU metrics (cf. M1,2) and the battery level (cf. M5,6) as critical, since system failures are imminent. For example, an increase of the CPU usage and temperature can lead from device shut down to overheating and core damage. Accordingly, the memory usage (cf. M3) is urgent, since it describes a state of heavy load, in which a decreasing free memory deteriorates the user experience (e. g., from slow apps to app shutdown). We classify the metrics with a better capacity to usage ratio like hard drive usage (cf. M4; high capacity) and the mobile data state (cf. M0) as intervention or recoverable, since they might lead to abnormal or unwanted behavior (e. g., “no disk space” or “expensive mobile data tariff”), however, can be mostly mitigated by the user, e. g., by adding an external SD-storage unit. The causal dependencies between the metrics are: $M1 \uparrow \rightarrow M2 \uparrow, M5 \downarrow$ (reads: increasing CPU usage leads to increasing temperature and battery drain); $M3 \uparrow \rightarrow M4 \downarrow, M4 \uparrow \rightarrow M3 \downarrow$ (data can be active in main memory or stored to disk).

2.2 Countermeasures

In our recent work on integration patterns we collected 22 distinct exception handling patterns (e. g., from prevention to fault tolerance) [17] as well as 7 (fault tolerant) integration adapter patterns [15,16]. Thereof eight patterns are applicable as countermeasures for the aforementioned impact metrics to add a resource-aware resilience, which we subsequently discuss. The other 21 patterns are not applicable, since they are rather focused on the message processing (i. e., dead letter channel, message validator [17]), or are abstract patterns (e. g., compensation or exception spheres [17]).

Fault Tolerance: Retry from Disk. The tolerance patterns: delayed retry, pause operation and request caching (all from [17]), can be combined to react on the mobile device being offline or using mobile data (cf. M0) by storing the monitoring (and the processed messages) to disk (countermeasures **CM-1: “store to disk”, CM-2: “stop storing to disk”**), instead of sending them to a remote endpoint. As soon as the situation allows (e. g., WIFI available), the current and the stored data can be delivered (**CM-3: “redeliver from disk, clean disk”**), and thus denoting three new resilience patterns.

Prevention: Throttle and Sample. A message throttler pattern [17] reduces the average message throughput of the system (**CM-4: “throttle processing”**)

and thus, reduces the CPU and battery usage (cf. M1,2,5,6) as well as reduces the data volume transferred per time unit (cf. M0). Since throttling reduces the messages by keeping them in main memory and emits only a certain ratio per time, the memory load might increase (cf. M3). This can be mitigated by using the message sampler pattern [17] by removing messages according to a pattern (**CM-5: “sample messages”**). This violates many service qualities (e. g., data loss), and thus the sampler should be used in critical or urgent situations only. A special case of the sampler in the context of “store to disk” (CM-1) is a parameterized sampling on existing (already stored data), which we call data store retention (**CM-6: “Store Retention”**), denoting another new pattern. In contrast to the sampler, the retention deletes messages on disk to insert new messages according to a policy (e. g., sample outdated messages on disk).

Escalation: Alert and Stop. In more severe situations, it might be necessary to either stop a part of the system (e. g., monitoring, integration scenarios) like the stop local pattern [17] (**CM-7: “stop processing”**) or even to stop the complete app [17] (**stop all; CM-8: “stop app”**). Since this leads to planned downtimes of parts or the complete system, an alert should precede the shut down to allow for a user intervention. The alert has no direct impact on the system, however, could be used to treat intervention and recoverable situations brought to the user in local or cloud monitors (**CM-9: “alert user”**).

CM Interdependencies. Although most of the countermeasures do not have side-effects, there are some that have to be discussed. The “stop app” (CM-8) impacts all other actions by stopping them together with the app. This does not further interfere with the other countermeasures after restart, since the situation will be newly evaluated and for example, the stored messages will be redelivered (CM-3). Furthermore, sample (CM-5) and redeliver from disk executed at the same time will only deliver parts of the stored data, which is acceptable for discrete resource and context data, variable with respect to its resolution (e. g., CPU, GPS coordinates), as common for context-aware apps.

3 Mobile Resilient Monitoring

In this section, we more formally define the behavior of our resilient EAI system monitor based on the countermeasures for resources and context from Sect. 2. This denotes the conceptual basis for a resource-aware, Event-Condition-Action [5] approach that executes countermeasures according to the state of the system.

3.1 Countermeasure Lifecycle: Resource and Network States

The allowed states of a resilient mobile system can be derived from the context situations found on mobile devices and the respective countermeasures (cf. both from Sect. 2) and represented as deterministic finite automata (DFA). Thereby we differentiate between a resource and a network automaton. The resource state

transitions are shown in Fig. 2(a), which denotes a DFA with the countermeasures (cf. CM-x) as states and the resource metric transitions as events (e.g., CPU temperature). When the app starts, the system is initially in normal processing, which changes as soon as memory, CPU or battery warning thresholds are crossed. Thereby, battery and CPU warnings include usage and temperature events. Whenever a new state is reached through an event transition, the countermeasure is executed by the system. Since this influences the overall system, the normal processing state can be reached from any countermeasure, as soon as all metrics transition into a normal state. The app can be shut down from any state and by external events. Common countermeasures to keep the system in allowed states are CM-4: “throttle processing”, CM-5: “sample messages” and CM-7: “stop processing”. When service qualities are about to degraded by the system (e.g., sample), CM-9: “alert user” is triggered to allow for interventions. Similarly, the network states are represented in Fig. 2(b). Notably, mostly CM-1–3 (“store to disk”, “stop storing to disk”, “redeliver from disk, clean disk”) are required. Instead of sampling, the data store retention (CM-6: “Store Retention”) is applied to free space for more recent messages.




Fig. 2. Context situations and countermeasures

3.2 Resilient State Representation


The event calculus (EC) [9] is a logic-based language for the formal specification of constraints, suitable for the reasoning on evolving traces, such as mobile context events. As for business monitoring approaches (e.g., [13]), the EC is applicable to our case, since situational changes can be represented over time, which allows to appropriately react on frequent, temporally grouped events (e.g., battery warning has to hold for a certain amount of time, before action applied).

EC Theories and Event Traces. The EC describes domain-independent predicates shown in Fig. 3(a) that allow for reasoning with first-order logic with negation on a (time-based) state transition system as in Fig. 2. For reasoning about time intervals we additionally define *clipped*(t_1, f, t_2), representing a fluent f (i.e., anything with a value subject to change over time) that is terminated between times t_1 and t_2 according to [21].

To more formally represent the countermeasure lifecycle and detect actionable system states we use EC theories similar to business monitoring [13], for

$\text{happens}(E, T)$	Event E happens at time T
$\text{holds_at}(F, T)$	Fluent F holds at time T
$\text{initially}(F)$	Fluent F holds in the initial state of the system
$\text{initiates}(E, F, T)$	Event E initiates fluent F at time T
$\text{terminates}(E, F, T)$	Event E terminates fluent F at time T

(a) Basic event calculus predicates from [9]



(b) Reactive reasoning, similar to [13]

Fig. 3. Event calculus predicates and reasoning

which we add mobile monitoring domain-specific fluents (e. g., countermeasures). Therefore we map all DFA states from Fig. 2 to a fluent (e. g., store) and transitions to an event (e. g., $e(-, \text{cpu_warn})$). Of these fluents, only one holds at a point in time, and the events are distinct with respect to a partial mobile event trace \mathfrak{T} (i. e., traces could extend in the future with new event occurrences), which is a finite set $\mathfrak{T} = \{\text{happens}(e(\text{id}, \text{type}()), t)\}$, where id is the event identifier, type is a function that combines the metrics predicates M1–M7 from Sect. 2 with a severity label to events, e. g., $\{\text{cpu_warn}(), \text{battery_limit}(), \dots\} \cup \{\text{app_starting}(), \dots\}$, and t is a timestamp. Further we add auxiliary fluents a_x for each event type (e. g., a_{bat} for bat_warn) to memorize, which event lead to the current state. Events are fired, e. g., as $\text{happens}(e(-, \text{cpu_warn}()), t)$, when the assigned predicates hold, e. g., $\text{cpu_warning}(\text{cpu}_{temp}, \text{cpu}_u) :- \text{cpu}(\text{cpu}_{temp}, \text{cpu}_u), \geq(\text{cpu}_u, th)$ with configurable threshold th .

R1: Device online to offline
 $\text{initially}_P(a_{net}). \text{initially}_P(a_{disk_free}).$

$\text{holds_at}(\text{norm}, t) \Rightarrow \text{initiates}(e(-, \text{no_net}()), \text{store}, t) \wedge$
 $\text{term.}(e(-, \text{no_net}()), \text{norm}, t) \wedge \text{term.}(e(-, \text{no_net}()), a_{net}, t).$
 $\text{holds_at}(\text{store}, t) \Rightarrow \text{term.}(e(-, \text{msgrcv}()), a_{disk_free}, t).$

R2: Device offline to online
 $\text{holds_at}(\text{store}, t) \wedge \text{holds_at}(a_{disk_free}, t) \Rightarrow$
 $\text{init.}(e(-, \text{net}()), \text{norm}, t) \wedge \text{term.}(e(-, \text{net}()), \text{store}, t).$
 $\text{holds_at}(\text{store}, t) \wedge \neg \text{holds_at}(a_{disk_free}, t) \Rightarrow$
 $\text{init.}(e(-, \text{net}()), \text{retry}, t) \wedge \text{term.}(e(-, \text{net}()), \text{store}, t).$
 $\text{holds_at}(\text{retry}, t) \Rightarrow \text{init.}(e(-, \text{msgsent}()), a_{disk_free}, t).$
 $\text{holds_at}(\text{retry}, t) \wedge \text{holds_at}(a_{disk_free}, t) \Rightarrow$
 $\text{init.}(e(-, \text{net}()), \text{norm}, t) \wedge \text{term.}(e(-, \text{net}()), \text{retry}, t).$

R3: Battery, CPU temp. \rightarrow CPU warning
 $\text{initially}_N(a_{bat}). \text{initially}_N(a_{cpu}).$

$\text{holds_at}(\text{norm}, t) \wedge \neg \text{holds_at}(a_{bat}, t) \Rightarrow$
 $\text{init.}(e(-, \text{batwarn}()), a_{bat}, t).$
 $\text{holds_at}(\text{norm}, t) \wedge \neg \text{holds_at}(a_{cpu}, t) \Rightarrow$
 $\text{init.}(e(-, \text{cpuwarn}()), a_{cpu}, t).$

$\text{holds_at}(a_{bat}, t) \wedge \neg \text{clipped}(t - 3, a_{bat}, t) \Rightarrow$
 $\text{term.}(e(-, \text{cpuwarn}()), \text{norm}, t) \wedge$
 $\text{init.}(e(-, \text{cpuwarn}()), \text{throttle}, t).$
 $\text{holds_at}(a_{cpu}, t) \wedge \neg \text{clipped}(t - 3, a_{cpu}, t) \Rightarrow$
 $\text{term.}(e(-, \text{cpuwarn}()), \text{norm}, t) \wedge$
 $\text{init.}(e(-, \text{cpuwarn}()), \text{throttle}, t).$

Example. The above table briefly shows excerpts of the overall theories for three cases that we require in the evaluation Sect. 5: rules R1–R3. The system is in a normal state at startup time t_1 , if $\text{holds_at}(\text{norm}, t_1)$. Most notably, a single event, e. g., $\text{happens}(e(-, \text{cpu_warn}()), t_2) \in \mathfrak{T}$, does not change the system's state. In contrast to R1 and R2, the system in R3 moves into another state, only if this fluent holds for the last three time units $t - 3$ with $\neg \text{clipped}(t - 3, a_{cpu}, t)$, and thus prevents from frequent state changes.

We define an implicitly included app_stopped fluent, which marks the termination of the app and all its integration processes when $\text{holds_at}(\text{app_stopped}, t)$, since a trace can be stopped at any time (not shown).

3.3 Efficient Reasoning

For the execution of the countermeasures, we reason about the defined EC theories and fire, if the system is in an actionable state (e.g., *throttle*, *retry*). Adapted to the resource-constraint devices, the reasoning should require minimal resources. While the common abductive reasoning (i.e., tries to generate a trace) is not applicable to our case, the deductive or backward reasoning (i.e., continuously evaluates conjunctions of holds_at predicates) is computationally expensive [21], and thus problematic on mobile devices. Hence we apply a reactive reasoning from [4], where the fluents’ validity intervals are revised or extended, when new events occur. This requires a cached EC (CEC) [4], which caches the maximal validity interval (MVI), in which a fluent holds uninterruptedly. Upon a new event occurrence CEC uses the cached MVIs to compute the new result. We adopt an approach similar to CEC shown in Fig.3(b), with a given initial state (initially), a mobile event trace stream (happens) and app event effects (initiates or terminates). The reasoner evaluates the MVIs of the evolving fluents (i.e., holds_at) and allows for the execution of countermeasures on an event stream with a time-based window of five minutes.

4 A Resilient, Mobile Integration System

For the evaluation, we briefly describe a realization of a resilient, resource-aware mobile integration system as well as the required extensions compared to existing EAI architectures [11]. Figure 4 shows the components of a conventional system with consumer and producer adapters communicating with sender and receiver applications as well as the integration process engine that executes the integration scenarios. Therefore, (operational) data and security-relevant stores are required. The optional message queues offer support for additional service qualities like flow control, or are used as message buffer.

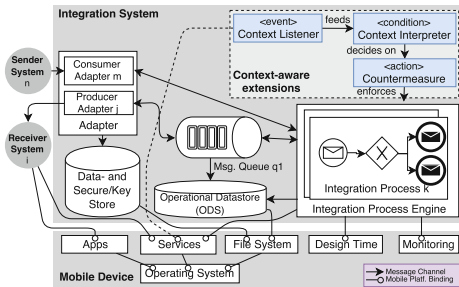


Fig. 4. Mobile EAI architecture.

enable the access to mobile services and other apps from integration processes. For the resilience, i.e., event-condition-action concept from Sect.3, we require

Architecture Extensions. Mobile devices have an operating system with app frameworks as well as services and file system access (incl. external storage). Although native OS realizations are possible, we assume the mobile integration system itself to be an app, which allows for more rapid prototyping (e.g., implement the stores and queues on these layers and write to the file system). In addition to transport protocols (e.g., HTTP, UDP), adapters

three new, vital components: context listener (collecting events: happens), context interpreter (evaluates conditions: holds_at), and a countermeasure (executor). While the listener—bound to the OS services—receives a stream of resource and context events, it converts them to EC theories according to Sect. 3, which are fed into the interpreter, evaluating the series of events and reacting by a list of countermeasures that are applied by the executor. Thereby the monitoring sub-system is represented by integration processes in a separate integration process instance, making the resilience characteristics applicable for non-monitoring integration scenarios.






Fig. 5. Camel app.

Realization. With more than two billion apps and a market share of approximately 90%², the predominant mobile OS is Android with Java as its major app programming language. Consequently, we ported the open source integration system Apache Camel version 2.17.0 [7] to Android as app  and added our monitor . Since the used EC approach relies on Horn clause fragments of first-order logic (with negation), we extended the system by our 150 kB low footprint datalog system [18], to address the context-aware extensions and the reasoning parts from Sect. 3. All experiments are conducted on a Samsung Galaxy S7 Edge with 2.3 Ghz QuadCore and 1.6 Ghz QuadCore ARM Cortex-A53, 4 GB main memory, 32 GB disk and 16

GB microSD, and Android Nougat OS. While the end-to-end communication tests via WIFI use Apache JMeter, the monitoring cases are covered by a based cloud monitor³. Figure 5 shows integration scenarios in form of Camel routes defined on the mobile device, where each route can be configured locally and local services and apps (e. g., OBD-II reader) can be used as adapters. While the system is running, all monitoring and resilience features can be switched-on and off by the user.

5 Evaluation

In this section, we evaluate our approach realized as mobile app (cf. Sect. 4) according to the following hypotheses, mirroring research question Q4:

- H1 The monitor approach does not have a significant impact on the message processing (e. g., \mapsto no increase of processing latency or memory consumption),
- H2 The service qualities are preserved (e. g., \mapsto avoid data loss in external monitor, up to a certain point: cf. trade-off resources vs. qualities).

Message Processing with Monitoring and Resilience (cf. H1). First we study the relation between the **normal** message processing of our mobile EAI

² Statista—Global mobile OS market share, visit. 02/2018: <http://bit.ly/2d7iCPb>.

³ Grafana, InfluxDB, visit. 02/2018: <http://docs.grafana.org/features/influxdb/>.

system and its execution with **monitoring** and **resilience**. Only when monitoring is switched on, we assume that the mobile event trace is collected for visualization according to Fig. 3(b), while the countermeasures are determined and executed only if the resilience is enabled. Figure 6(a) denotes the results of the end-to-end latency of messages of different sizes sent from Apache JMeter to a mobile Apache Camel route over WIFI until the response is received. Although the WIFI was exclusively used during the benchmark without any other networks in reach, the mean of the five runs slightly varies for the measurement in the 16kB message size case. For the other message sizes, the results show no significant deterioration of the latency with monitoring or resilience compared to the normal processing (i. e., no monitoring and resilience) with a 0.95 confidence interval (cf. error bars). Similarly, there is no significant impact on the memory consumption with increasing messages sizes between the normal processing and our extensions as shown in Fig. 6(b). We indicate this through the secondary y-axis, which shows the consumption difference $\text{diff} = \max(\text{mem}_{\text{normal}}, \text{mem}_{\text{monitor}}, \text{mem}_{\text{resilience}}) - \text{mem}_{\text{normal}}$, with a relatively stable diff even with a growing amount of data during message processing. The memory consumption of the monitoring and resilience remain stable, even for an increasing number of senders (not shown).




Fig. 6. Side-effects of monitoring and resilience on normal EAI processing.

Trade-Off: Limited Resources vs. No Data Loss (cf. H2). We consider no data loss as a service quality during online-offline switches, which can happen quite frequently for a mobile device. Therefore, we switched-off the WIFI on the device at time t_1 , while trying to deliver the monitoring data to a remote cloud app (cf. Sect. 4), which we used as a remote monitor. Consequently the cloud monitor does not receive any further data points (cf. Fig. 7(a) WIFI Signal), which are written to the device’s disk according to rule R-1 (implementing CM-1) from $t_1 \approx 3:20\text{pm}$ UTC onwards, and thereby increasing the local disk usage (cf. HDD usage). When we switched-on the WIFI at time t_i ($i > 1$), rule R-2 stops the storage to disk and starts the redelivery, which filled the data points in the cloud monitor (i. e., no gaps in Fig. 7(a)) and freed disk resources (cf. CM-2,3).



Fig. 7. Resilient monitoring in Grafana cloud dashboard.

Battery. Furthermore, the battery and CPU temperatures are considered critical service qualities, for which we show the effect of an increasing message load in the integration system that let the CPU and **Battery temperature** rise together with an increased CPU usage as shown in Fig. 7(b). When the countermeasure (e. g., throttling) is executed according to rule R-3, then the effect is mitigated without data loss (cf. from $\approx 2:35$ am UTC). Since the temperature reduces below the threshold, the state switches from throttling to normal processing.

Discussion. Briefly, we experimentally showed that the normal EAI processing is invariant to the new EAI system extensions proposed in this work (cf. H1). Furthermore, we illustrated two cases, in which service qualities were preserved: no data loss and battery temperature. However, with respect to the trade-off: with longer offline executions (cf. Fig. 7(a)), leading to an increased disk usage, data loss cannot be prevented beyond the disk limit. Since the monitoring and resilience sub-systems are constructed by using Apache Camel routes, and thus integration processes, the results of this work mainly discussed for the monitoring case, are directly applicable to the actual EAI runtimes (e. g., #threads, #msgs/sec).

6 Related Work

Although there is a growing market for mobile automation software (e. g., IFTTT or Integromat Mobile), we are not aware of any work on resource-aware, resilient EAI systems on mobile devices. This observation is supported by recent work on mobile process [20], service engines [19] for data collection scenarios and surveys in the field [1], especially mentioning the lack of a resource-aware monitoring and energy efficiency of current solutions. For a more systematic analysis, we conducted a literature review following [8] across disciplines with the combinations of the keywords: android, application, context-aware, device, mobile, monitoring, resources, resource-aware, rule-based, quality of service (from [15]). The horizontal search resulted in 63 articles out of which we selected nine, few from the years between 2004–2005, but mostly 2012–2016, indicating recent interest in this topic.

Service Engines. According to the analysis, most of the related work is from the SOA domain, designing mobile service engines comparable to [19]. The closest known related work from [3] defines a dual system resource approach (i. e.,

remote and local service). As the work from [25], the approach is limited to network-awareness (i. e., minimize communication) and as countermeasure they shut down the network (cf. M0, Sect. 2). The early work from [6] describes another mobile SOA monitoring approach of an embedded system in a car. However, none of them describe a mobile EAI system and they do not monitor the execution on the mobile system, and thus do not describe a resilient approach. They were not evaluated from a functional runtime (e. g., latency) and resilience perspective.

Miscellaneous. The work from [2, 23], mainly focus on mobile network observation and can react to issues with the network (cf. M0, Sect. 2). [12] defines a central monitoring of a distributed system, by generating the network of monitoring trees that optimizes multiple monitoring tasks and balances the resource consumption at different nodes. Although the focus is different (e. g., no resource-restrictions) the work can be seen complementary to ours. Similar to our approach, [27] introduces a monitoring system that is evaluated on an android device. However, the monitoring focuses exclusively on the user's health in form of a fitness app. The work by [26] uses a rule-based approach to monitor a process system on android similar to [20]. Again, both approaches do not specify resource-constraints or countermeasures. [14] defines a rule-based mobile phone system, however, for context-aware, mobile applications.

7 Discussion

Toward a resilient mobile application integration and monitoring, we proposed solutions along the research questions Q1–4, by (a) an analysis of data, impact factors and countermeasures, (b) a formal mobile monitoring foundation based on event calculus that allows for the necessary reasoning over the changing system state for a resilient behavior of the monitoring as well as the runtime, (c) an architecture extension for a resource-aware EAI system, and (d) evaluated the approach with a mobile app prototype. Although this work mainly defines the contributions (a)–(c) around a resilient monitoring use-case, the result is applicable to integration runtime processing as well. Since the integration scenarios can become complex, future work should address on-device modeling. Furthermore, the presented approach requires the definition of thresholds, whose definition for all metrics could be difficult and might require adaptations along changing system characteristics. This opens another research direction toward parametrized countermeasures with constraints for automatic system adaptation and optimization, as already indicated by CM-6.

Acknowledgments. We thank Jan Lentmaier for the implementation support.

References

1. Aceto, G., Botta, A., De Donato, W., Pescapè, A.: Cloud monitoring: a survey. *Comput. Netw.* **57**(9), 2093–2115 (2013)
2. Arora, H., Greenwald, L.: Toward the use of local monitoring and network-wide correction to achieve QoS guarantees in mobile ad hoc networks. In: *IEEE SECON*, pp. 128–138 (2004)
3. Chen, S., et al.: Context-aware resource management middleware for service oriented applications based on cross-layer monitoring. In: *IEEE SMC* (2015). Page submitted
4. Chittaro, L., Montanari, A.: Efficient temporal reasoning in the cached event calculus. *Comput. Intell.* **12**(3), 359–382 (1996)
5. Dittrich, K.R., et al.: The active database management system manifesto: a rule-base of ADBMS features. In: Sellis, T. (ed.) *RIDS*, pp. 1–17 (1995)
6. Gehlen, G., Mavromatis, G.: A rule based data monitoring middleware for mobile applications. In: *IEEE VTC*, vol. 5, pp. 2910–2914 (2005)
7. Ibsen, C., Anstey, J.: *Camel in Action*. Manning Publications Co., Greenwich (2010)
8. Kitchenham, B.: *Procedures for Performing Systematic Reviews*, vol. 33 (2004)
9. Kowalski, R., Sergot, M.: A logic-based calculus of events. In: Schmidt, J.W., Thanos, C. (eds.) *Foundations of Knowledge Base Management. Topics in Information Systems*, pp. 23–55. Springer, Heidelberg (1989). https://doi.org/10.1007/978-3-642-83397-7_2
10. Ligu, S.: *Effective Monitoring and Alerting*. O'Reilly Media Inc. (2012)
11. Linthicum, D.S.: *Enterprise Application Integration*. Addison-Wesley, Boston (2000)
12. Meng, S., Kashyap, S.R., Venkatramani, C., Liu, L.: Resource-aware application state monitoring. *IEEE TPDS* **23**(12), 2315–2329 (2012)
13. Montali, M., Maggi, F.M., Chesani, F., Mello, P., van der Aalst, W.M.: Monitoring business constraints with the event calculus. *ACM TIST* **5**(1), 17 (2013)
14. Nalepa, G.J., Bobek, S.: Rule-based solution for context-aware reasoning on mobile devices. *Comput. Sci. Inf. Syst.* **11**(1), 171–193 (2014)
15. Ritter, D., Holzleitner, M.: Integration adapter modeling. In: *CAiSE*, pp. 468–482 (2015)
16. Ritter, D., May, N., Rinderle-Ma, S.: Patterns for emerging application integration scenarios: a survey. *Inf. Syst.* **67**, 36–57 (2017)
17. Ritter, D., Sosulski, J.: Exception handling in message-based integration systems and modeling using BPMN. *IJCIS* **25**(02), 1650004 (2016)
18. Ritter, D., Westmann, T.: Business network reconstruction using datalog. In: Barceló, P., Pichler, R. (eds.) *Datalog 2.0 2012. LNCS*, vol. 7494, pp. 148–152. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-32925-8_15
19. Schobel, J., Pryss, R., Wipp, W., Schickler, M., Reichert, M.: A mobile service engine enabling complex data collection applications. In: Sheng, Q.Z., Stroulia, E., Tata, S., Bhiri, S. (eds.) *ICSOC 2016. LNCS*, vol. 9936, pp. 626–633. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46295-0_42
20. Schobel, J., Pryss, R., Schickler, M., Reichert, M.: A lightweight process engine for enabling advanced mobile applications. In: Debruyne, C., et al.: (eds.) *On the Move to Meaningful Internet Systems: OTM 2016 Conferences. OTM 2016. LNCS*, vol. 10033, pp. 552–569. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-48472-3_33

21. Shanahan, M.: The event calculus explained. In: Wooldridge, M.J., Veloso, M. (eds.) *Artificial Intelligence Today*. LNCS (LNAI), vol. 1600, pp. 409–430. Springer, Heidelberg (1999). https://doi.org/10.1007/3-540-48317-9_17
22. Smanchat, S., Ling, S., Indrawan, M.: A survey on context-aware workflow adaptations. In: *MoMM*, pp. 414–417. ACM (2008)
23. Tiemeni, G.L.N.: Performance estimation of wireless networks using traffic generation and monitoring on a mobile device. Ph.D. thesis (2015)
24. Vukolić, M.: The Byzantine empire in the intercloud. *ACM SIGACT News* **41**(3), 105–111 (2010)
25. Wagh, K.S., Thool, R.C.: Monitoring web resources and improve QoS of mobile host using fuzzy logic. **4**, 1779–1784 (2014)
26. Wipp, W.: *Workflows on Android: A Framework Supporting Business Process Execution and Rule-Based Analysis*. Ph.D. thesis, Ulm University (2016)
27. Zhang, W., Wang, X.: A lightweight user state monitoring system on android smartphones. In: Toumani, F., et al. (eds.) *ICSOC 2014*. LNCS, vol. 8954, pp. 259–269. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-22885-3_23

Social Media and Web-Based Business Information Systems



Tight and Loose Coupling in Evolving Platform Ecosystems: The Cases of Airbnb and Uber

Andreas Hein , Markus Böhm , and Helmut Krcmar 

Technical University of Munich, Munich, Germany
{andreas.hein, markus.boehm, krcmar}@in.tum.de

Abstract. The emergence of digital platforms changes the way how companies interact with their ecosystem. Successful platforms like Apple’s App Store utilize an ecosystem of third-party developers to drive innovation. Those platforms are expanding the sphere of influence beyond internal resources and capabilities by taking advantage of a scalable ecosystem of external developers. However, until now it is unclear on how those companies establish a platform ecosystem. This article draws on two case studies in the form of ridesharing and accommodation platforms to illustrate how they transitioned through four evolutionary phases with the help of tight and loose coupling partnerships.

Keywords: Digital platform · Platform ecosystem · Ecosystem evolution

1 Introduction

The success of digital platforms is unheard [1]. Examples range from application platforms like Apple’s App or Google’s Play Store to the currently emerging IoT platform market. In either case, the surrounding ecosystem of partners, developers, and users strongly determines the success of the platform. In contrast to traditional businesses, where the company produces products or services, a platform orchestrates the interaction in a two-sided market of suppliers (e.g., App developers) and users (e.g., App users) [2–4]. The ecosystem fosters cross-side network effects between both sides, where each side profits from additional exchange [2, 3]. By capturing a margin of the created value, the platform has a self-interest in growing this ecosystem. Current practices in starting an ecosystem that incorporates tight coupling of selected parties through individual partnership to a scalable utilization of boundary resources like Application Programming Interfaces (APIs) and Software Development Kits (SDKs) resulting in a loosely coupled partnerships [5, 6]. Based on those practices, ecosystems go through several phases starting with the birth, expansion, leadership, and self-renewal [7]. Those rapidly emerging digital platform ecosystems pose challenges to traditional companies. They face the question of how to start an ecosystem effectively to remain competitive.

This article sheds light on this question by elaborating on the ecosystem evolution of the recent phenomenon of transactional service platforms like Uber and Airbnb. Both provide a servitization in the form of mobility or accommodation services based on software artifacts. Uber, for example, faced little competition and managed to grow an

ecosystem of developers by providing APIs and SDKs resulting in new possibilities to support their core mobility service. The ridesharing company uses the ecosystem to co-create new software services to enhance, for example, the trip experience of passengers. Airbnb, on the other hand, managed to survive in a shark tank of competition. The platform became one of the highest evaluated start-ups by growing an ecosystem that surpassed the competition by fostering tight coupling partnerships. Also, Airbnb established securing mechanisms to maintain authority and control. Airbnb finally opened their ecosystem in late 2017 for external developers. Both cases provide valuable insights on the use of different coupling mechanisms during the ecosystem evolution to shed light how to start and grow a platform ecosystem.

2 Theoretical Background

This article draws on the theoretical groundwork of Moore [7] on business ecosystems and the theory of Weick [8] on tightly and loosely coupled systems. We combine both lenses to elaborate on how platform ecosystems evolved.

2.1 The Platform Ecosystem

The theory of business ecosystem is shaped by influences from the field of Anthropology and co-evolution [9], as well as from Biology and natural ecosystems [10]. Co-evolution emphasizes that interdependent species or actors evolve in nearly endless reciprocal cycles. The evolution of one species, thus, influences the natural selection of another species in the same ecosystem and vice versa. On the other hand, natural ecosystems respond with sensitivity to environmental change [9, 10]. Moore combines both aspects to describe how businesses evolve under ever-changing environmental influences in an ecosystem of partners, suppliers, and customers. The theory proposes four evolutionary stages [7]:

Birth: The first evolutionary phase describes risk-taking entrepreneurs that focus on customer needs and how to satisfy those needs. An example is the rise of the personal computer in the 1970s, which started in 1975 by the sole effort of hobbyists in hard-to-use products. Only in the late 1970s, companies like Apple utilized an ecosystem of business partners in the field of independent software developers, training institutes or computer stores. In this ecosystem, Apple tightly controlled the computer design and the operating software, while encouraging and co-evolving with other partners to contribute complementary software or hardware [7].

The first phase harbors the cooperative challenge of collaborating with the demand and supply side to create new value propositions that are the basis for future innovations. At the same time, the young ecosystem needs to protect its value proposition from imitation by tightly integrating the parties involved in the ecosystem [7].

Expansion: In the second phase, business ecosystems expand to new territories. This phase is characterized by direct battles for market share between competing or overlapping ecosystems. Companies that aim to expand their ecosystem need to have a

scalable business concept in place. An example for capturing territory can be found in the personal computer ecosystem when IBM entered and replaced Apple's dominant position. IBM opened its architecture to hardware supplier fostering additional complementary products. Also, they licensed MS-DOS and ensured the compatibility and portability of popular applications like Wordstar [7].

The cooperative challenge is to establish a scalable ecosystem of supply and demand to achieve maximum market coverage. From a competitive perspective, the company needs to establish a de-facto industry standard by dominating key market segments [7].

Leadership: The ecosystem leadership phase expresses stable structures and a high concentration of bargaining power by a dominating company. The leader creates products or services that are of critical value for the whole ecosystem and protects this position through patents or constant innovation. The central ecological position of the ecosystem leader is often instantiated by clearly defined interfaces to partners reflecting the industry standard. Though those interfaces, the ecosystem leaders encourage partners in the ecosystem to take over activities and to accelerate the growth of the whole. In the case of the personal computer ecosystem, missing leadership and a too open architecture opened a window of opportunity for competitors like Lotus, Intel, and Microsoft. Microsoft and Intel proofed to exercise control over critical components in the form of the operating systems and chips needed for the personal computer [7].

The cooperative challenges imply to provide a compelling vision for the future that encourages the ecosystem to work together to strengthen the central position of the leader. On the other hand, the competitive challenge is to maintain central bargaining power in the ecosystem [7].

Self-Renewal: The last phase of business ecosystems covers self-renewal or death. Established or incumbent ecosystems are threatened by risk-taking actors and innovations. A typical example of this stage is Apple's comeback with the iPod and later iPhone software ecosystem. The company managed to overhaul incumbent ecosystems like Nokia's Symbian operating system. Thus, leading and sustaining an ecosystem through continually adapting and absorbing technological advances is crucial to long-term success. Countermeasures of platform leaders are slowing the growth or enveloping competing ecosystems [2, 7].

The cooperative challenges in this phase is that ecosystem leaders need to work closely with innovators to incorporate new ideas into their ecosystem. From a competitive perspective, the ecosystem leader should focus on maintaining high entrance barriers to prevent competitors from invading and switching costs for users to prevent them from leaving the ecosystem [7].

2.2 Tight and Loose Coupling

Tight and loose coupling describes the degree of dependency between actors within and between organizations or systems. Orton and Weick [6] differentiate between tight, loose, and decoupled systems:

Tight Coupling: In a tightly coupled system, elements are strongly dependent on each other, and they do not act independently. Researchers speak of responsiveness without distinctiveness [6, 8]. Typical characteristics of tightly coupled systems are an increased understanding of each other's needs, a close relationship, a low degree of information asymmetry, and the ability to tailor products or services to strategic needs [8, 11]. A practical example for tight coupling is a strategic partnership like seen in the IT outsourcing, where the partnership is determined by precise rules like Service-Level-Agreements and a high degree of mutual dependency [12]. **Loose Coupling:** The concept of loose coupling refers to independent elements that are distinct or separate from one another, yet responsive. Scientists call this order distinctive responsiveness. The advantage of a system is that elements maintain flexibility to react to change from outside but adhere to standards in the system. Thus, the system maintains stability as change from external sources cause no ripple effects. As a result, the system is simultaneously closed and open, flexible and stable, as well as rational and spontaneous [6]. **Decoupling:** In opposition to neither distinctive nor responsive elements, which is per definition no system, decoupled elements are distinct to each other [6].

3 Research Approach

The research design follows a multiple case study with subject to accommodation and ridesharing platforms. The method is particularly suitable as it captures and describes the complexity of real events [13]. The cases show boundaries, features, and limitations, by putting the business ecosystems [7] and tight/loose coupling [6, 8] theories into a specific context including the respective environment and firms [14]. Benbasat, Goldstein and Mead [15] provide an orientation on whether the analyzed topic is a phenomenon and the usage of a case study is appropriate. First, the context is crucial to observe the phenomenon of emerging platform ecosystems opposed to an isolated view on platforms. Secondly, the tremendous success of platform businesses like Airbnb and Uber shows the significance and actuality of the research topic and link to the contemporary event of emerging platform ecosystems [13]. Additionally, neither control nor manipulation of the subject or event took place. This was guaranteed as the case study describes the phenomenon in the view of a neutral observer [16]. After conducting the two case studies, we draw cross-case conclusions on how a platform ecosystem evolves by showing similarities and differences between the cases.

The data for the two cases are based on archival data and empirical studies, ranging from the emergence of the respective platform ecosystem to the present. Archival data includes financial data, independent reports on milestones, partnerships, key events (like introducing an API), and the general development of the platform and its ecosystem. We used additional empirical studies to validate financial figures, events, as well as strategic decisions. Further, we adhered to the principle of data triangulation to make sure that each fact or decision is cross-validated [17]. Another mechanism to take care of business model specifics was including the main competitor as a benchmark in the ecosystem. Overall, we collected 24 archival data sources and 12 empirical studies on

the case of accommodation platform ecosystems, as well as 28 archival data sources and 16 empirical studies on the ridesharing platform ecosystem case (see Table 1).

Table 1. Data used for the multiple-case study

Industry	Case	Archival data ^a	Empirical studies
Accommodation	AirBnb	14	12
	Competitors (HomeAway)	10	
Ridesharing	Uber	18	16
	Competitors (Lyft)	10	

^a Archival data ranges from financial reports found on: Forbes, Bloomberg, VentureBeat, CB Insights, Crunchbase; and independent reports from: TechCrunch, Programmableweb, Business Insider, Fortune, CNBC, The Guardian, The Verge, Forbes, and company websites.

4 Results

The results show that both platform ecosystems used tight coupling to evolve from birth to the expansion phase. In a second stage, both adhered to a more open approach through loose coupling resulting in platform leadership and self-renewal.

Birth: Founded in 2009, Uber triggered the trend of on-demand ridesharing companies. Instead of competing directly with the taxi ecosystem but rather with the concept of owning a car, Uber managed to collect \$ 1.75 million during their first two rounds of seed investment in 2009 and 2010. The company started in the San Francisco Bay Area and provided “high-end” sedans with reputable drivers that can be easily booked via iPhone taking advantage of the GPS sensors. Uber did not want to compete for the lowest price but charged 1.5x of the regular taxi rate offering a comfortable and convenient way of moving from A to B. Another distinguishing feature is the use of state-of-the-art technology incorporating the iPhone sensors to provide an easy to use experience. Uber managed to establish a position in the mobility ecosystem by matching the demand of convenient mobility services with the supply of professional chauffeurs.

Most people tend to think that Airbnb had always been the number one platform when it comes to short-term lodging. The truth is that Airbnb was surrounded by other vacation rental companies like HomeAway or Flip-key founded in 2005 and 2007 from the very beginning. In 2008, the company called AirBed & Breakfast entered a highly competitive ecosystem offering short-term living quarters in the center of San Francisco. During the initial phase, Airbnb managed to keep up with its competitors by providing additional value-adding services resulting from tight-coupled partnerships. One of them was proposing free and professional photography from the housings on the platform to reduce, on the one hand, the effort for the host to join the platform and to increase, on the other hand, the reliability and trustworthiness for the guest. At the end of 2011, Airbnb established its position in the ecosystem and raised over \$ 1 billion in funding, which was nearly half the amount of HomeAway with 2.2 billion dollars.

Expansion: After introducing the iOS and Android Apps in 2010, Uber launched a national expansion starting in the San Francisco Bay Area, expanding to New York City in 2011. Until the year 2012, Uber collected over \$ 1.9 billion in funding and could develop their business without any direct competition in the ecosystem. In the same year, Lyft emerged as the first direct competitor. Lyft dates back to Zimride, which was a long-distance ridesharing company. Lyft has the same focus as its competitor with on-demand ridesharing in a peer-to-peer manner. Ending in the year 2012, Uber used the three years without competition to establish its position with a valuation of \$ 1.9 billion compared to \$ 275 million of Lyft. Based on this leading edge, Uber established various tightly coupled partnerships to secure and expand on its current position in the ecosystem. One exemplary partnership that strengthened the supply side (drivers) was providing special interest rates on car loans for GM and Toyota vehicles [18]. Newly bought cars that engage in the Uber platform can be seen as assets to lower the interest rates for potential drivers ensuring scalability on the supply side. Other partnerships involve PayPal and American Express to offer additional payment services reducing the entrance barriers for passengers or the corporation with Concur for business travelers or flight procedures within United Airlines to expand the market reach [19]. With the help of those tightly managed partnerships, Uber could extend its position in the ecosystem and establish itself as a de-facto standard until 2014 with a valuation of \$ 41 billion compared to 2.5 \$ billions of Lyft.

Airbnb needed to be more careful as they acted in a highly competitive ecosystem. One example is that Airbnb offered an active API endpoint including an affiliation program for developers [20]. However, new competition in the form of Wimdu and 9Flats emerged. Airbnb could not establish enough authority and consequently did not officially introduce the programming interface. Besides closing the API end-point, Airbnb also launched a new flagging feature to defend its content by reporting questionable behavior. The response is somewhat understandable, as both new competitors were known to copy established business models including their content [21]. Airbnb focused on additional tight coupling partnerships to grow the ecosystem and to progress further. Examples are the partnerships with Lloyds of London to offer insurances for hosts included in every platform booking or the affiliation with American Express introducing new payment functionalities [22]. All those partnerships aim to improve the relationship between the host and the guest, in turn, affecting the reputation of Airbnb. At the end of the year 2015, Airbnb harvested the fruits of their expansion with \$ 25.5 billion followed by HomeAway with an evaluation of \$ 3.9 billion. Besides, also the number of listings (Airbnb: from 50,000 listings in 2011 to 550,000 in 2014) and active drivers (Uber: from near zero in 2012 to 160,000 in 2014) rose sharply [23, 24].

Leadership and Self-Renewal: Starting the next phase, Uber established enough authority in the ecosystem to respond to the demand of third-party developers by opening up the platform through Application Programming Interfaces (APIs) and Software Development Kits (SDKs). Starting in 2014, third-party developers could integrate core mobility services provided by Uber, like ordering a ride with the help of predefined code snippets. Further, Uber opened the micro-service architecture of the platform to help developers using driver, car, and ride data to develop new applications, and to enhance

the driving experience during the trip [25]. One example of tweaking the open mobility services is the use of calling an Uber via an SMS over the phone. Those apps emerge out of the loosely coupled relationship between developers in the ecosystem and are especially useful in areas without mobile internet. Uber is not only using the ecosystem to expand the market reach but also takes advantage of the innovation capabilities of external developers. They control the ecosystem through the strategic use of components in the form of APIs like the Trip Experience API or the UberRUSH API strengthening their core service. One example is the integration of contextual information during the ride that Uber provides to third-party developers. They use the information to create value-adding applications that benefit both Uber and the passenger. Consequently, Lyft followed this strategy and introduced an open API in 2016. As of 2017, Uber has established a vibrant ecosystem of developers providing five different APIs, 22 SDKs in several programming languages and three software libraries [26]. They transitioned from a tight coupling approach to target special needs like supporting potential drivers owning a car to a more loosely coupled approach to utilize a whole ecosystem of developers to foster constant innovation.

In the accommodation industry, it was after 2015 and onward, as Airbnb's competitor HomeAway first opened their platform through active API endpoints including documentation, guides, and an affiliate program. The market need for the content data of those accommodation platforms was omnipresent. In this period, several companies tried to scrape the data from the Airbnb platform, like AirDNA analyzing the content of Airbnb and selling it for investment purposes or several attempts to reverse engineer the not yet publicly open Airbnb API endpoints. All those efforts point toward the direction that Airbnb controlled the core components of the platform ecosystem. Finally, the combination of all those market signals and the economic advance and authority regarding market valuation to the leading competitors influenced Airbnb's decision towards a more open and loosely coupled approach. Arrived at the year 2017, Airbnb emerged as a winner of the shark tank of short-term lodging platforms and finally taking control over their ecosystem by introducing an official API program including documentation and partnership programs.

5 Discussion

The two cases show that platform ecosystems follow the evolutionary pattern of business ecosystems. In addition, the theoretical lens of tight and loose coupling helps to explain how platforms exercise scalability through network effects to drive innovation. The results indicate that the platforms use tight coupling partnerships in the birth and expansion phase to become ecosystem leaders, followed by a transition towards openness and loosely coupled partnerships to drive external innovation (see Fig. 1). Instead of dominating each step in the value chain as, for example, Original Equipment Manufacturers do in the automotive industry; platform ecosystems are in their mature states more loosely coupled and open, which fosters open innovation effects.

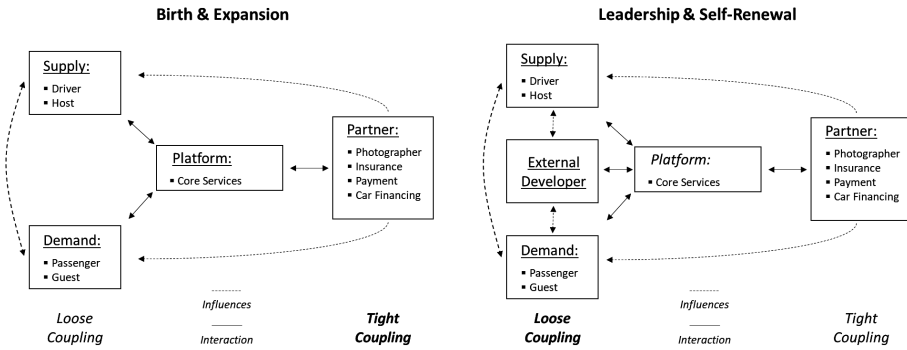


Fig. 1. The use of tight and loose coupling during the evolution of platform ecosystems

5.1 Tight Coupling: Controlling the Ecosystem

During the birth phase, both platforms exercised a high degree of control over their ecosystem. Airbnb needed to deal with competition and utilized tight coupling partnerships to increase the scope of the platform, thus, maximizing the value for customers. An example is the collaboration with professional photographers to increase the trust between guests and hosts. Uber had no direct competition in the first place but managed to differentiate its value proposition through a more convenient and reliably way of inner-city traveling. The phase of expansion dealt with an emerging competition and encouraged securing and expanding mechanisms. The appearance of additional competition forced Airbnb to retain control over their content by closing all open API endpoints. Also, the platform established tight coupling partnerships like insuring apartments with companies to further increase the trust between hosts and guests. On the other hand, Uber used tight coupling partnerships on the supply side by collaborating with GM to provide special interest rates for potential drivers out-competing other platforms. Both, GM and Uber developed mutual trust, where Uber convinced GM that drivers could use the car as a profit generation resource, while GM convinced Uber to provide safe and reliable vehicles. Another example is the cooperation with Concur to expand to business travelers. The left part of Fig. 1 illustrates how both platforms took advantage of tight coupling partnerships to indirectly influence the value proposition between the demand and supply side. Airbnb offered a professional photo shooting for free, creating value for the host through the polished appearance of the room, as well as value for the guest as all pictures made by Airbnb were verified reducing the risk of being tricked. The same is true for the special car financing loans for Uber drivers. Passengers enjoy brand-new cars and drivers have the opportunity to earn an independent living. The platforms indirectly enhanced cross-side network effects between the supply and demand-side through strategic or tightly coupled partnerships. The aim of those partnerships is either to strengthen the value proposition between the two sides as seen in the photograph or car financing example or to expand the ecosystem as with Concur. Tight coupling partnerships lead to a high degree of control over the strategic

intention and output of the collaboration. In return, both parties need to allocate resources to align on a mutual strategy and outcomes.

5.2 Loose Coupling: Opening the Ecosystem

When eventually opening the platform for loose coupling, both platforms achieved leadership in their respective ecosystems surpassing the competition regarding market valuation with a factor of 16 in the case of Uber vs. Lyft and 8 in the case of Airbnb vs. HomeAway. The two platforms managed to obtain authority over core services and strong bargaining power, which provided the opportunity to open the ecosystem through APIs and SDKs to directly induce network effects and open innovation. Airbnb achieved this state just recently in end 2017 when announcing an official API due to the ongoing competition. Uber opened the platform already in 2014 providing until now five different APIs and 22 SDKs to their ecosystem of developers. External developers can use boundary resources to co-create additional value around the core mobility service. Each API, follows a predefined interface structure in the form of data payloads but also provides flexibility for the developer to use the service however they want. Based on this loose coupling partnership, the external developers can use their unique background to satisfy special (long-tail) needs. A similar pattern can be found in the app store ecosystems of Apple and Google. Figure 1 shows how ecosystem leadership enabled platforms to open their ecosystems to benefit from external innovations. Instead of tightly integrating and optimizing new value-adding services around their core components (e.g., mobility or accommodation matching), platforms engage their ecosystem to constantly innovate on the platform. The arms-length cooperation reduces the intensity of the partnership in exchange for lower coordination efforts and control over the ecosystem. This loose coupling allows the platforms to untap the innovation potential of a rapidly growing ecosystem.

5.3 How to Open the Platform?

According to Uber, the primary intention behind opening the platform was based on the three pillars of utility, revenue, and distribution [27]. Uber sees utility as a continuous improvement process of making the platform profit from new ideas within the ecosystem. They utilize a compelling vision for the future of mobility-as-a-service going hand in hand with Moore's description of ecosystem leadership. The second pillar aims to incentivize the development of new applications in the form of affiliate programs that encourage ecosystem partners to contribute toward core components. Lastly, the distribution pillar provides value during additional distribution (e.g., improving the experience during the ride). Uber provides stability by leading the co-innovation through targeted APIs like the Trip Experience API enabling developers to embed complementary content, games or other functionality that enhance the experience during trips. As Airbnb just opened their platform, it is at a too early stage to predict what complementary services emerge around the ecosystem. However, the early approaches of setting up an affiliate program point in a somewhat similar direction. As a prerequisite for opening the platform, both platforms established control and authority over their ecosystem to

deal with competition. The examples of Uber and Airbnb show the utilization of tight coupling partnerships to build the core ecosystem. Both platforms used value-adding partnerships to enhance the scope and range of the platforms core services, as well as securing mechanisms to stay on the leading edge regarding competition [5]. After expanding to a dominant position, both platforms opened their ecosystem around crucial platform services that evolved toward a market standard. They integrated external developers in loosely coupled partnerships through boundary resources like APIs and SDKs. Other platforms from industries like Social Networks and App Stores follow a similar approach. From the majority of loosely coupled developers' point of view, it makes sense to invest time and resources in an already established market standard, instead of multi-homing between several competing platforms with an increased effort. A technical commonality is that both companies transitioned from a monolithic to a modular micro-service infrastructure that can be adapted to key products or services. With this approach, the platform makes sure that newly created products extend the main product or service. Uber, for example, opened the ecosystem for third party development to add additional value to passengers during the ride fostering complementary innovation. To achieve this goal, both platforms use scalable interfaces between the modular micro-service infrastructure and the ecosystem. SDKs, libraries, and documentation help to simplify the development process by covering the most used programming languages to include pre-defined functions and snippets helping developers to get started. By transitioning from a tightly coupled ecosystem to a more arms-length approach, the platform governance needs to be adapted as well as shown in the example of the incentive programs. Other aspects of the ecosystem range from coping with the on- and off-boarding of developers in a scalable way to the automatic input and output control of applications within the ecosystems [28, 29].

6 Conclusion, Implications, and Future Research

This research provides insights into evolutionary phases and coupling mechanisms in the context of platform ecosystems. We show that platforms follow the evolutionary phases as proposed by Moore [7]. However, there are differences when introducing the mechanisms of tight and loose coupling. The two cases indicate that platforms use tight coupling partnerships in the beginning to enhance the core service or value proposition strengthening cross-side network effects between the supply and demand-side. Thus, tight coupling is one promising mechanism to answer the open question on how platforms strategize to foster cross-side network effects [30]. Another theoretical contribution points towards the increasing role of value co-creation in a platform ecosystem [31]. We highlight that platforms first prepare the ecosystem and then enable third-party developers to actively innovate around core services. For practitioners, we show that timing is crucial when it comes to the question when to open a platform. Airbnb learned that authority and control over their ecosystem were important and that the platform's content can be scraped by competitors if not protected properly. Thus, it was important to first strengthen core services through strategically selected tight coupling partnerships and then opening up to utilize the innovation capability of a loosely coupled ecosystem.

In this sense, a fruitful area for future research could be the identification of value co-creation practices with emphasize on the tangible character of tight coupling partnerships and loosely-coupled or intangible ecosystem relationships. Also, new technology like blockchain might influence the future development of platform ecosystems. The research also faces limitations. The data for the two case studies are mainly based on archival data and empirical research taking the role of the neutral observer. We thus limit the explanatory power of this contribution on answering the motives or exact practices of both platforms when using tight and loose coupling mechanisms. One way to mitigate this problem is conducting in-depth case studies to elaborate on the why and how of tight and loose coupling practices. In total, this research helps to explain how platforms accelerate network effects in their ecosystem evolution. They first establish ecosystem leadership with the help of strategical or tight coupling partnerships that strengthen the value proposition of their core service. Then, they unleash the network externalities by opening the platform to the ecosystem and enabling third parties in loosely coupled partnerships through clearly defined interfaces to innovate around those core services.

Acknowledgements. This work is part of the TUM Living Lab Connected Mobility (TUM LLCM) project and has been funded by the Bavarian Ministry of Economic Affairs and Media, Energy and Technology (StMWi) through the Center Digitisation. Bavaria (ZD.B), an initiative of the Bavarian State Government.

References

1. Evans, P.C., Gawer, A.: The rise of the platform enterprise: a global survey (2016)
2. Eisenmann, T., Parker, G., Alstyne, V.M.W.: Strategies for two-sided markets. *Harvard Bus. Rev.* **84**, 92–101 (2006)
3. Rochet, J.C., Tirole, J.: Platform competition in two-sided markets. *J. Eur. Econ. Assoc.* **1**, 990–1029 (2003)
4. Hein, A., Böhm, M., Krcmar, H.: Platform Configurations within Information Systems Research: A Literature Review on the Example of IoT Platforms, pp. 465–476. Multikonferenz Wirtschaftsinformatik, Lüneburg, Germany (2018)
5. Ghazawneh, A., Henfridsson, O.: Balancing platform control and external contribution in third-party development: the boundary resources model. *Inf. Syst. J.* **23**, 173–192 (2013)
6. Orton, J.D., Weick, K.E.: Loosely coupled systems: a reconceptualization. *Acad. Manag. Rev.* **15**, 203–223 (1990)
7. Moore, J.F.: *The Death of Competition: Leadership and Strategy in the Age of Business Ecosystems*. Harper Business, New York (1997)
8. Weick, K.E.: Educational organizations as loosely coupled systems. *Adm. Sci. Q.* **21**, 1–19 (1976)
9. Bateson, G.: *Mind and Nature: A Necessary Unity*. Dutton, New York (1979)
10. Gould, S.J.: *The Structure of Evolutionary Theory*. Harvard University Press (2002)
11. Danneels, E.: Tight–loose coupling with customers: the enactment of customer orientation. *Strateg. Manag. J.* **24**, 559–576 (2003)
12. Steensma, H.K., Corley, K.G.: On the performance of technology-sourcing partnerships: the interaction between partner interdependence and technology attributes. *Acad. Manag. J.* **43**, 1045–1067 (2000)

13. Yin, R.K.: *Case Study Research: Design and Methods*. Sage Publications, Thousand Oaks (2014)
14. Merriam, S.B.: *Case Study Research in Education: A Qualitative Approach*. Jossey-Bass, San Francisco (1988)
15. Benbasat, I., Goldstein, D.K., Mead, M.: The case research strategy in studies of information systems. *MIS Q.* **11**, 369–386 (1987)
16. Siggelkow, N.: Persuasion with case studies. *Acad. Manag. J.* **50**, 20–24 (2007)
17. Klein, H.K., Myers, M.D.: A set of principles for conducting and evaluating interpretive field studies in information systems. *MIS Q.* **23**, 67–93 (1999)
18. Tam, D.: Uber teams up with GM, Toyota to get more cars on the road, C-net. <https://www.cnet.com/news/uber-teams-up-with-gm-toyota-to-get-more-cars-on-the-road/>. Accessed 01 Nov 2017
19. Hempel, J.: Uber and the future of American Express, *Fortune*. <http://fortune.com/2014/06/09/uber-and-the-future-of-american-express/>. Accessed 1 Nov 2017
20. DuVander, A.: AirBNB API Will Pay Developers with Affiliate Program, ProgrammableWeb. <https://www.programmableweb.com/news/airbnb-api-will-pay-developers-affiliate-program/2011/10/13>. Accessed 01 Feb 2018
21. Tsotsis, A.: Airbnb Freaks Out over Samwer Clones, *TechCrunch*. <https://techcrunch.com/2011/06/09/airbnb/>. Accessed 1 Nov 2017
22. Empson, R.: Airbnb Makes Good, Will Now Cover Up To \$1M in Property Damages, *TechCrunch*. <https://techcrunch.com/2012/05/22/airbnb-1m-guarantee/>. Accessed 01 Nov 2017
23. Hall, J.V., Krueger, A.B.: An analysis of the labor market for Uber’s driver-partners in the United States. *ILR Review* 0019793917717222 (2015)
24. Ferenstein, G.: Uber and Airbnb’s incredible growth in 4 charts, *VentureBeat*. <https://venturebeat.com/2014/06/19/uber-and-airbnbs-incredible-growth-in-4-charts/>. Accessed 01 Feb 2018
25. Haddad, E.: Service-Oriented Architecture: Scaling the Uber Engineering Codebase as We Grow, *Uber Engineering*. <https://eng.uber.com/soa/>. Accessed 01 Nov 2017
26. Programmable Web: Uber API, *Programmable Web*. <https://www.programmableweb.com/api/uber>. Accessed 1 Nov 2017
27. Wintrob, G.: Talking to Uber about the three pillars of its API platform, *TechCrunch*. <https://techcrunch.com/2016/06/16/talking-to-uber-the-three-pillars-of-ubers-api-platform/>. Accessed 01 Feb 2018
28. Hein, A., Schrieck, M., Wiesche, M., Krcmar, H.: Multiple-case analysis on governance mechanisms of multi-sided platforms. In: *Multikonferenz Wirtschaftsinformatik, Ilmenau, Germany*, pp. 1613–1624 (2016)
29. Schrieck, M., Hein, A., Wiesche, M., Krcmar, H.: The challenge of governing digital platform ecosystems. In: *Linnhoff-Popien, C., Schneider, R., Zaddach, M. (eds.) Digital Marketplaces Unleashed*, pp. 527–538. Springer, Heidelberg (2018). https://doi.org/10.1007/978-3-662-49275-8_47
30. McIntyre, D.P., Srinivasan, A.: Networks, platforms, and strategy: emerging views and next steps. *Strateg. Manag. J.* **38**, 141–160 (2017)
31. Lusch, R.F., Nambisan, S.: Service innovation: a service-dominant logic perspective. *MIS Q.* **39**, 155–175 (2015)



On Feeding Business Systems with Linked Resources from the Web of Data

Andrea Cimmino^(✉) and Rafael Corchuelo^(✉)

ETSI Informática, University of Seville,
Avda. Reina Mercedes, s/n., Sevilla, Spain
{cimmino, corchu}@us.es

Abstract. Business systems that are fed with data from the Web of Data require transparent interoperability. The Linked Data principles establish that different resources that represent the same real-world entities must be linked for such purpose. Link rules are paramount to transparent interoperability since they produce the links between resources. State-of-the-art link rules are learnt by genetic programming and build on comparing the values of the attributes of the resources. Unfortunately, this approach falls short in cases in which resources have similar values for their attributes, but represent different real-world entities. In this paper, we present a proposal that leverages a genetic programming that learns link rules and an ad-hoc filtering technique that boosts them to decide whether the links that they produce must be selected or not. Our analysis of the literature reveals that our approach is novel and our experimental analysis confirms that it helps improve the F_1 score by increasing precision without a significant penalty on recall.

1 Introduction

The feasibility of many emerging business relies on the availability and interoperability of suitable on-line datasets [1]. The Web of Data provides business with islands of data (availability) that can be transparently used as required by the business models (interoperability). The Web of Data is ruled by the Linked-Data principles that support the idea that resources within different datasets that represent the same real-world entities must be linked so as to facilitate data interoperability [4]. Link rules are intended to help link resources automatically however they may be linked by different kind of relations, e.g., spatial or temporal coverage, we focus only on owl:sameAs that relates resources representing the same real-world entity.

Supported by the Spanish R&D programme (grants TIN2013-40848-R and TIN2016-75394-R). The computing facilities were provided by the Andalusian Scientific Computing Centre (CICA). We are grateful to Dr. Carlos R. Rivero and Dr. David Ruiz for earlier ideas that led to the results in this paper. We also thank Dr. Francisco Herrera for his hints on statistical analyses and sharing his software with us.

The literature provides several proposals to machine learn link rules by means of genetic programming [11, 12, 19, 20]. Such rules build on transformation and similarity functions that are applied to the values of the attributes of two resources to check if they can be considered similar enough (by attributes we mean their data-type properties); if they are, then the input resources are linked; otherwise, they are kept apart. Such link rules fall short where there are many resources that represent different real-world entities but have attributes with similar values.

In this paper, we present a novel approach to the problem: first, we leverage a state-of-the-art genetic programming to learn a set of link rules; we then select a link rule and apply it in order to obtain a collection of candidate links; the remaining rules are then boosted to analyse the neighbours of the resources involved in each candidate link (the neighbours are the resources that can be reached by means of their object properties); finally, we analyse how similar they are in order to decide which of the candidate links must be selected as true positives and which must be discarded as false positives. Our analysis of the related work unveils that this is a novel approach since current state-of-the-art link rules do not take the neighbours into account. Our experimental analysis confirms that precision can be improved by 77% in average, with an average -9% impact on recall; overall, the average improvement regarding the F_1 score is 58%. We also conducted the Iman-Davenport test to check that these differences are statistically significant regarding precision and the F_1 score, but not regarding recall. Our conclusion is that ours is a good approach to help software agents integrate the data that they fetch from the Web of Data, which will definitely help many business systems have access to many datasets in the Web of Data and integrate them with others.

A hypothetical business that may benefit from our approach is a collector that aims at selling food products from several supermarkets choosing always the one with the lower fair. In this scenario it is paramount to identify the products that are the same in the different supermarkets, so the cheapest can be suggested. Furthermore the collector may also consider other datasets to add business value to their data. For instance a food-health dataset and a datasets containing recipes, ingredients, menus and diets which along the food-health data may allow to sell packages of food that conform healthy menus and recipes. As a result, relying on a highly precise approach like ours to identify the same products sparse along the different datasets is paramount.

The rest of the article is organised as follows: Sect. 2 reports on the related work; Sect. 3 provides the details of our proposal; Sect. 4 presents our evaluation results; finally, Sect. 5 summarises our conclusions.

2 Related Work

The earliest techniques to learn link rules were devised in the field of traditional databases, namely: de-duplication [7, 17], collective matching [2, 3, 6, 14, 21], and entity matching [15]. They set a foundation for the researchers who addressed

the problem in the context of the Web of Data, where data models are much richer and looser than in traditional databases.

Some of the proposals that are specifically-tailored to web data work on a single dataset [10, 16], which hinders their general applicability; there are a few that attempt to find links between different datasets [5, 8, 9, 13], but they do not take the neighbours of the resources being linked into account, only the values of the attributes; that is, they cannot make resources with similar values for their attributes apart in cases in which they represent different real-world entities. An additional problem is that they all assume that data are modelled by means of OWL ontologies. Unfortunately, many common datasets in the Web of Data do not rely on OWL ontologies, but on simple RDF vocabularies that consists of classes and properties whose relationships and constraints are not made explicit.

The previous problems motivated several authors to work on techniques that are specifically tailored to work with RDF datasets. Most such proposals rely on genetic programming [11, 12, 19, 20] in which chromosomes encode the link rules as trees, which facilitates performing cross-overs and mutations. They differ regarding the expressivity of the language used to encode the link rules and the heuristics used to implement the selection, replacement, cross-over, and mutation operators, as well as the performance measure on which the fitness function relies. Isele and Bizer [11, 12] contributed with a supervised proposal called Genlink. It is available with the Silk framework [24], which is gaining impetus thanks to many real-world projects [23]. It uses a tournament selection operator, a generational replacement operator, custom cross-over and mutation operators, and its fitness function relies on the Matthews correlation coefficient. It can use a variety of custom string transformation functions and the Levenshtein, Jaccard, Numeric, Geographic, and Date string similarity measures. An interesting feature is that the size of the link rules must not be pre-established at design time, but it is dynamically adjusted during the learning process. Ngomo and Lyko [19] contributed with a supervised proposal called Eagle, which is available with the LIMES framework [18]. It uses a tournament selection operator, a $\mu + \beta$ replacement operator, tree cross-over and mutation operators, and its fitness function relies on the F_1 score. It does not use transformation functions, but the Levenshtein, Jaccard, Cosine, Q-Grams, Overlap, and Trigrams string similarity functions. The maximum size of the link rules must be pre-established at design time. Nikolov et al. [20] contributed with an unsupervised proposal. It uses a roulette-wheel selection operator, an elitist replacement operator, a tree cross-over operator, a custom mutation operator, and a pseudo F_1 fitness function. Transformations are not taken into account, but the library of similarity functions includes Jaro, Levenshtein, and I-Sub. The maximum size of the link rules rules is also set at design time. There is a diverging proposal by Soru and Ngomo [22]. It supports the idea of using common machine-learning techniques on a training set that consists in a vectorisation of the Cartesian product of the resources in terms of the similarity of their attributes. Transformation functions cannot be used and the string similarity functions available are Q-Grams, Cosine, and Levenshtein. Whether the size of the rules must be pre-set or not depends on the underlying machine learning technique.

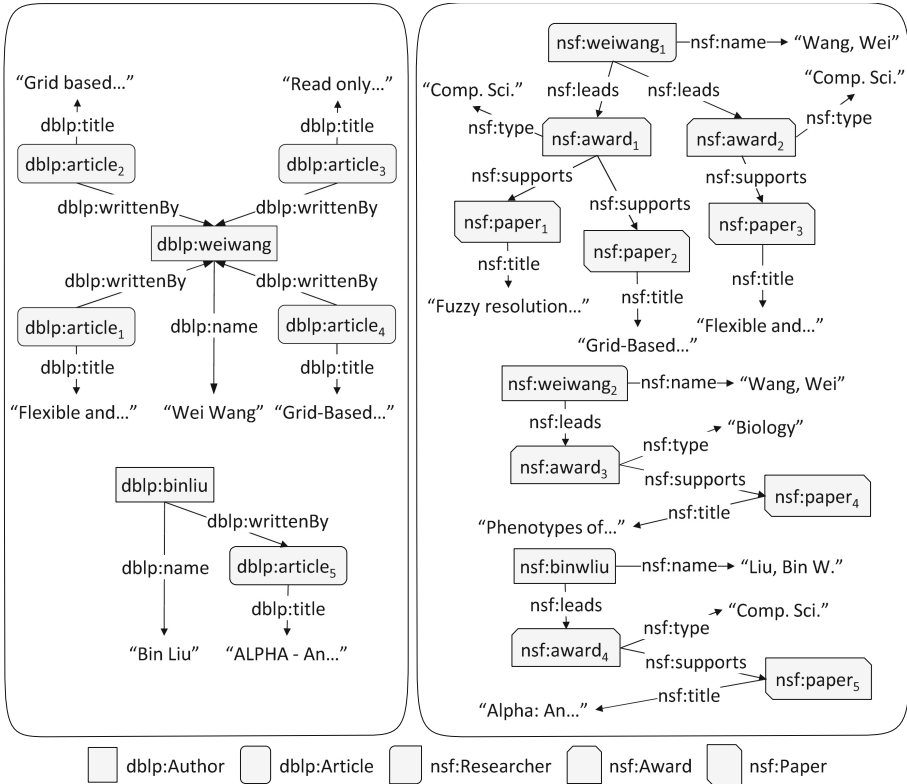


Fig. 1. Running example.

Unfortunately, none of the proposals for RDF datasets take the neighbourhoods into account (Fig. 1).

Clearly, the state of the art does not account for a proposal to link resources in RDF datasets that takes their neighbours into account. Ours is tailored to work with such datasets and it is novel in that it is not intended to generate link rules, but leverages the rules that are learnt with other proposals and boosts them in order to analyse the neighbours, which our experimental analysis confirms that has a positive impact on precision without degrading recall.

3 Our Proposal

Our proposal consists in three components, namely: the first one learns link rules, the second one filters out the links that they produce, and the third relies on two voting strategies to select most reliable filtered links.

The link rule learner is based on Genlink [12], which is a state-of-the-art genetic programming that has proven to be able to learn good link rules for many common datasets. Genlink is feed with a set of examples representing

resources that should be linked, then it learns a set of link rules from them. As a result Genlink returns a set of link rules (all of them aim at linking the same kind of resources). A quality score ranks the output rules, usually the best is selected and the rest discarded. Doing this task several times for different kind of resources a set of link rules is generated, one of these rules is selected to be improved using our approach and the rest are the supporting rules.

The filter is an ad-hoc component that works as follows: it takes a link rule and executes it to produce a set of candidate links; then, it analyses the neighbours of the resources involved in each candidate link by boosting the remaining rules; links in which the corresponding neighbours are similar enough are preserved as true positive links while the others are discarded as false positive links. The selector is an ad-hoc component that works as follows: it takes the filtered links and the supporting rules used to filter them; then, it performs a voting strategy regarding how many rules filtered the same link and another voting strategy regarding how many links were filtered by each rule; finally, a subset of the filtered links is selected and preserved as the rest are discarded.

Below, we present the details of the filter and selector, plus an ancillary method that helps measure how similar the neighbours of two resources are.

Example 1. Listing 1 presents two sample datasets that are based on the DBLP and the NSF datasets. The resources are depicted in greyed boxes whose shapes encode their classes (i.e., the value of property *rdf:type*), the properties are represented as labelled arrows, and the literals are encoded as strings. The genetic component learns the following link rules in this scenario, which we represent using a Prolog-like notation for the sake of readability:

$$\begin{aligned}
 r_1: \text{link}(A, R) \text{ if } & \text{rdf:type}(A) = \text{dblp:Author}, \text{rdf:type}(R) = \text{nsf:Researcher}, \\
 & N_A = \text{dblp:name}(A), N_R = \text{nsf:name}(R), \\
 & \text{levenstein}(\text{lfname}(N_A), \text{lfname}(N_R)) > 0.80. \\
 \\
 r_2: \text{link}(A, P) \text{ if } & \text{rdf:type}(A) = \text{dblp:Article}, \text{rdf:type}(P) = \text{nsf:Paper}, \\
 & T_A = \text{dblp:title}(A), T_P = \text{nsf:title}(P), \\
 & \text{jaccard}(\text{lowercase}(T_A), \text{lowercase}(T_P)) > 0.65.
 \end{aligned}$$

where *levenstein* and *jaccard* denote the well-known string similarity functions (normalised to interval [0.00, 1.00]), *lfname* is a function that normalises people's names as "last name, first name", and *lowercase* is a function that changes a string into lowercase.

Intuitively, link rule r_1 is applied to a resource A of type *dblp:Author* and a resource R of type *nsf:Researcher*; it computes the normalised Levenshtein similarity between the normalised names of the author and the researcher; if it is greater than 0.80, then the corresponding resources are linked. Link rule r_2 should now be easy to interpret: it is applied to a resource A of type *dblp:Article* and a resource P of type *nsf:Paper* and links them if the normalised Jaccard similarity amongst the lowercase version of the title of article A and the title of paper P is greater than 0.65.

```

1: method filterLinks( $r, S, D_1, D_2, \theta, \mu, \rho$ ) returns  $K$ 
2:    $K := \emptyset$ 
3:    $(C_1, C_2) := (\text{sourceClasses}(r), \text{targetClasses}(r))$ 
4:    $L_1 := \text{apply}(r, D_1, D_2)$ 
5:   for each link rule  $r' \in S$  do
6:      $(C'_1, C'_2) := (\text{sourceClasses}(r'), \text{targetClasses}(r'))$ 
7:      $(P_1, P_2) := (\text{findPath}(C_1, C'_1, D_1), \text{findPath}(C_2, C'_2, D_2))$ 
8:      $L_2 = \text{apply}(r', D_1, D_2)$ 
9:     for each  $(p_1, p_2) \in P_2 \times P_2$  do
10:      for each link  $(a, b) \in L_1$  do
11:         $(A, B) := (\text{findResources}(a, p_1, D_1), \text{findResources}(b, p_2, D_2))$ 
12:         $E := L_2 \cap (A \times B)$ 
13:         $w := \text{computeSimilarity}(A, B, E)$ 
14:        if  $w \geq \theta$  then
15:           $K := K \cup \{(a, b)\}$ 
16:        end
17:      end
18:    end
19:  end
20:   $K := \text{selectLinks}(K, \mu, \rho)$ 
21: end

```

Listing 1. Method to filter links.

It is not difficult to realise that link rule r_1 links resources *dblp:weiwang* and *nsf:weiwang₁* or *dblp:binliu* and *nsf:binwliu*, which are true positive links, but also *dblp:weiwang* and *nsf:weiwang₂*, which is a false positive link. In cases like this, the only way to make a difference between such resources is to analyse their neighbours, be them direct (e.g., *dblp:weiwang* and *dblp:article₂*) or transitive (e.g., *nsf:weiwang₁* and *nsf:paper₂*). ■

3.1 Filtering Links

Listing 1 presents the method to filter links. It works on a link rule r , a set of supporting link rules S , a source dataset D_1 , a companion dataset D_2 , and a threshold θ that we explain later. It returns K , which is the subset of links produced by base link rule r that seem to be true positive links.

The method first initialises K to an empty set, stores the source and the target classes of the base link rule in sets C_1 and C_2 , respectively, and the links that result from applying it to the source and the companion datasets in set L_1 .

The main loop then iterates over the set of supporting link rules using variable r' . In each iteration, it first computes the sets of source and target classes involved in link rule r' , which are stored in variables C'_1 and C'_2 , respectively; next, it finds the set of paths P_1 that connect the source classes in C_1 with the source classes in C'_1 in dataset D_1 ; similarly, it finds the set of paths P_2 that connect the target classes in C_2 with the target classes in C'_2 in dataset D_2 . By

path between two sets of classes, we mean a sequence of object properties that connect resources with the first set of classes to resources with the second set of classes, irrespective of their direction. Simply put: the idea is to find the way to connect the resources linked by the base link rule with the resources linked by the supporting link rule, which is done by the intermediate and the inner loops.

The intermediate loop iterates over the set of pairs of paths (p_1, p_2) from the Cartesian product of P_1 and P_2 . If there is at least a pair of such paths, it then means that the resources involved in the links returned by base link rule r might have some neighbours that might be linked by supporting link rule r' .

The inner loop iterates over the collection of links (a, b) in set L_1 . It first finds the set of resources A that are reachable from resource a using path p_1 in source dataset D_1 and the set of resources B that are reachable from resource b using path p_2 in the companion dataset D_2 . Next, the method applies supporting link rule r' to the source and the companion dataset and intersects the resulting links with $A \times B$ so as to keep resources that are not reachable from a or b apart; the result is stored in set E . It then computes the similarity of sets A and B ; intuitively, the higher the similarity, the more likely that resources a and b refer to the same real-world entity. If the similarity is equal or greater than threshold θ , then link (a, b) is added to set K ; otherwise, it is filtered out. When the main loop finishes, set K contains the collection of links that involve neighbours that are similar enough according to the supporting rules.

We do not provide any additional details regarding the algorithms to find paths or resources since they can be implemented using Dijkstra's algorithm to find the shortest paths in a graph. Computing the similarity coefficient is a bit more involved, so we devote a subsection to this ancillary method below.

Example 2. In our running example, link rule r_1 is the base link rule, i.e., we are interested in linking authors and researchers, and we use link rule r_2 as the support link rule, i.e., we take their articles and papers into account. Their source classes are $C_1 = \{dblp:Author\}$ and $C'_1 = \{dblp:Article\}$, respectively, and their target classes are $C_2 = \{nsf:Researcher\}$ and $C'_2 = \{nsf:Paper\}$, respectively. Link rule r_1 returns the following links: $L_1 = \{(dblp:weiwang, nsf:weiwang_1), (dblp:weiwang, nsf:weiwang_2), (dblp:binliu, nsf:binwliu)\}$; note that the first and the third links are true positive links, but the second one is a false positive link. Link rule r_2 returns the following links: $L_2 = \{(dblp:article_1, nsf:paper_3), (dblp:article_2, nsf:paper_2), (dblp:article_4, nsf:paper_2), (dblp:article_5, nsf:paper_5)\}$, which are true positive links.

The sets of paths between the source and target classes of r_1 and r_2 are $P_1 = \{\langle dblp:writtenBy \rangle\}$ and $P_2 = \{\langle nsf:leads, nsf:supports \rangle\}$. Now, the links in L_1 are scanned and the resources that can be reached from the resources involved in each link using the previous paths are fetched.

Link $l_1 = (dblp:weiwang, nsf:weiwang_1)$ is analysed first. The method finds $A = \{dblp:article_1, dblp:article_2, dblp:article_3, dblp:article_4\}$ by following resource $dblp:weiwang$ through path $\langle dblp:writtenBy \rangle$; similarly, it finds $B = \{nsf:paper_1, nsf:paper_2, nsf:paper_3\}$ by following resource $nsf:weiwang_1$ through path $\langle nsf:leads, nsf:supports \rangle$. Now supporting link rule r_2 is applied and the results

```

1: method computeSimilarity(A, B, E) returns d
2:   A' := reduce(A, E)
3:   B' := reduce(B, E)
4:   W := intersect(A', B', E)
5:   d := |W| / min{|A'|, |B'|}
6: end

```

Listing 2. Method to compute similarity.

are intersected with $A \times B$ so as to keep links that are related to l_1 only; the result is $E = \{(dblp:article_1, nsf:paper_3), (dblp:article_2, nsf:paper_2), (dblp:article_4, nsf:paper_2)\}$. Then, the similarity of A and B in the context of E is computed, which returns 0.67; intuitively, there are chances that l_1 is a true positive link.

Link $l_2 = (dblp:weiwang, nsf:weiwang_2)$ is analysed next. The method finds $A = \{dblp:article_1, dblp:article_2, dblp:article_3, dblp:article_4\}$ by following resource *dblp:weiwang* through path $\langle dblp:writtenBy \rangle$; next, it finds $B = \{nsf:paper_4\}$ by following resource *nsf:weiwang₂* through path $\langle nsf:leads, nsf:supports \rangle$. Now supporting link rule r_2 is applied and the result is intersected with $A \times B$, which results in $E = \emptyset$. In such a case the similarity is zero, which intuitively indicates that it is very likely that l_2 is a false positive link.

Link $l_3 = (dblp:binliu, nsf:binwliu)$ is analysed next. The method finds $A = \{dblp:article_5\}$ by following resource *dblp:binliu* through path $\langle dblp:writtenBy \rangle$; next, it finds $B = \{nsf:paper_5\}$ by following resource *nsf:binwliu* through path $\langle nsf:leads, nsf:supports \rangle$. Now supporting link rule r_2 is applied and the result is intersected with $A \times B$, which results in $E = \{(dblp:article_5, nsf:paper_5)\}$. The similarity is now 1.00, i.e., it is very likely that link l_3 is a true positive link.

Assuming that θ is set to, e.g., 0.50, the *filterLinks* method would return $K = \{(dblp:weiwang, nsf:weiwang_1), (dblp:binliu, nsf:binwliu)\}$. Note that the previous value of θ is intended for illustration purposes only because the running example must necessarily have very little data. ■

3.2 Computing Similarity

Listing 2 shows our method to compute similarities. Its input consists of sets A and B , which are two sets of resources, and E , which is a set of links between them. It returns the Szymkiewicz-Simpson overlapping coefficient, namely:

$$overlap(A, B) = \frac{|A \cap B|}{\min\{|A|, |B|\}}$$

The previous formula assumes that there is an implicit equality relation to compute $A \cap B$, $|A|$, or $|B|$. In our context, this relation must be inferred from the set of links E by means of Warshall's algorithm to compute the reflexive, commutative, transitive closure of relation E , which we denote as E^* .

The method to compute similarities relies on two ancillary functions, namely: *reduce*, which given a set of resources X and a set of links E returns a set whose

elements are subsets of X that are equal according to E^* , and *intersect*, which given two reduced sets of resources X and Y and a set of links E returns the intersection of X and Y according to E^* . Their definitions are as follows:

$$\begin{aligned} \text{reduce}(X, E) &= \{W \mid W \propto W \subseteq X \wedge W \times W \subseteq E^*\} \\ \text{intersect}(X, Y, E) &= \{W \mid W \propto W \subseteq X \wedge \exists W' : W' \subseteq Y \wedge W \times W' \in E^*\} \end{aligned}$$

where $X \propto \phi$ denotes the maximal set X that fulfils predicate ϕ , that is:

$$X \propto \phi \Leftrightarrow \phi(X) \wedge (\nexists X' : X \subseteq X' \wedge \phi(X'))$$

The method to compute similarities then works as follows: it first reduces the input sets of resources A and B according to the set of links E ; it then computes the intersection of both reduced sets; finally, it computes the similarity using Szymkiewicz-Simpson's formula on the reduced sets.

Example 3. Analysing link $l_1 = (dblp:weiwang, nsf:weiwang_1)$ results in sets $A = \{dblp:article_1, dblp:article_2, dblp:article_3, dblp:article_4\}$, $B = \{nsf:paper_1, nsf:paper_2, nsf:paper_3\}$, and $E = \{(dblp:article_1, nsf:paper_3), (dblp:article_2, nsf:paper_2), (dblp:article_4, nsf:paper_2)\}$. If E is interpreted as an equality relation by computing its reflexive, symmetric, transitive closure, then it is not difficult to realise that $dblp:article_2$ and $dblp:article_4$ can be considered equal, because $dblp:article_2$ is equal to $nsf:paper_2$ and $nsf:paper_2$ is equal to $dblp:article_4$. Thus, set A is reduced to $A' = \{\{dblp:article_1\}, \{dblp:article_2, dblp:article_4\}, \{dblp:article_3\}\}$ and set B is reduced to $B' = \{\{nsf:paper_1\}, \{nsf:paper_2\}, \{nsf:paper_3\}\}$. As a conclusion, $|A' \cap B'| = |\{\{dblp:article_1, nsf:paper_3\}, \{dblp:article_2, dblp:article_4, nsf:paper_2\}\}| = 2$, $|A'| = 3$, and $|B'| = 3$; so the similarity is 0.67.

When link $l_2 = (dblp:weiwang, nsf:weiwang_2)$ is analysed, $A = \{dblp:article_1, dblp:article_2, dblp:article_3, dblp:article_4\}$, $B = \{nsf:paper_4\}$, and $E = \emptyset$. Since the equality relation E^* is then empty, the similarity is zero because the intersection between the reductions of sets A and B is empty.

In the case of link $l_3 = (dblp:binliu, nsf:binliu)$, $A = \{dblp:article_5\}$, $B = \{nsf:paper_5\}$, and $E = \{(dblp:article_5, nsf:paper_5)\}$. As a conclusion, $|A' \cap B'| = |\{\{dblp:article_5, nsf:paper_5\}\}| = 1$, $|A'| = 1$, and $|B'| = 1$, where A' and B' denote, respectively, the reductions of sets A and B ; so the similarity is 1.00. ■

3.3 Selecting Links

Listing 3 shows the method to select the best links out of a set of correspondences. Its input consists of a set of correspondences K , a threshold μ to the minimum number of times that a link must have been selected by a supporting link rule so that it can be selected by this method (top links), and an additional threshold ρ to the minimum number of times that a supporting link rule must have selected a link so that links that have been selected by that link rule can be selected by this method (top link rules) even if they are not top links (Fig. 2).

This method relies on two ancillary functions, namely: *links*, which given a supporting link rule r' returns the set of links that it has selected, and *rules*,

```

1: method selectLinks( $K, \mu, \rho$ ) returns  $L$ 
2:    $M := \{(s, t) \mid \exists r' : (s, t, r') \in K\}$ 
3:    $S := \{r' \mid \exists s, t : (s, t, r') \in K\}$ 
4:    $g := \mu \max\{n \mid \exists a, b : (a, b) \in M \wedge n = |\text{rules}(a, b, K)|\}$ 
5:    $d := \rho \max\{m \mid \exists r' : r' \in S \wedge m = |\text{links}(r', K)|\}$ 
6:    $U := \{(a, b) \mid (a, b) \in M \wedge |\text{rules}(a, b, K)| \geq g\}$ 
7:    $V := \{r' \mid r' \in S \wedge |\text{links}(r', K)| \geq d\}$ 
8:    $L := \{(a, b) \mid (a, b) \in M \wedge ((a, b) \in U \wedge \text{rules}(a, b, K) \cap V \neq \emptyset)\}$ 
9: end

```

Listing 3. Method to select filtered links.

Scenario	Genlink							Our proposal							
	links	P	R	F ₁	θ	μ	ρ	links	P	R	F ₁	Δ Links	Δ P	Δ R	Δ F ₁
DBLP - NSF	127	0.25	0.97	0.40	0.01	0.02	0.74	33	0.97	0.97	0.97	-94	0.96	0.00	0.95
DBLP - DBLP	78,348	0.12	1.00	0.21	1.00	0.07	0.03	9,069	1.00	1.00	1.00	-69,279	1.00	0.00	1.00
BBC - DBpedia	525	0.85	1.00	0.92	0.10	0.00	0.03	461	0.96	1.00	0.98	-64	0.74	0.00	0.72
DBpedia - IMDB	118	0.27	0.55	0.36	0.60	0.00	0.03	41	0.68	0.48	0.57	-77	0.57	-0.12	0.32
RAE - Newcastle	404	0.22	0.82	0.35	0.30	0.13	0.03	68	0.72	0.45	0.56	-336	0.64	-0.45	0.32
Rest - Rest	103	0.90	0.83	0.87	0.10	0.08	0.58	96	0.97	0.83	0.89	-7	0.69	0.00	0.19
												Average Δ	0.77	-0.09	0.58
												Iman-Davenport's test	0.00	0.25	0.00

Fig. 2. Experimental results.

which given a link (a, b) returns the set of link rules that have selected it. The previous functions are formally defined as follows:

$$\begin{aligned} \text{links}(r', K) &= \{(a, b) \mid \exists r' : (a, b, r') \in K\} \\ \text{rules}(a, b, K) &= \{r' \mid \exists a, b : (a, b, r') \in K\} \end{aligned}$$

The method to select links first projects the set of correspondences K onto the set of links M and the set of supporting link rules S . It then computes g as a percentage, according to μ , of the maximum number of link rules that have selected each candidate link; it also computes d as a percentage, according to ρ , of the maximum number of links that a support link rule has selected as candidates. Next, it computes the set of top links U as the set of links in M that have been selected by at least g link rules; similarly, it computes the set of top link rules V as the set of link rules in S that have selected at least d links. The resulting set of links L is computed as the subset of links in M that are either top links or have been selected by top link rules.

4 Experimental Analysis

In this section, we first describe our experimental environment and then comment on our results.

Computing facility: We run our experiments on a virtual computer that was equipped with four Intel Xeon E5-2690 cores at 2.60 GHz and 4 GiB of RAM. The operating system was CentOS Linux 7.3.

Prototype:¹ We implemented our proposal with Java 1.8 and the following components: the Genlink implementation from the Silk Framework 2.6.0 to generate link rules, Jena TDB 3.2.0 to work with RDF data, ARQ 3.2.0 to work with SPARQL queries, and Simmetrics 1.6.2, SecondString 2013-05-02, and JavaStringSimilarity 1.0.1 to compute string similarities.

Evaluation datasets:² We used the following datasets: DBLP, NSF, BBC, DBpedia, IMDb, RAE, Newcastle, and Rest. We set up the following scenarios: (1) DBLP–NSF, which focuses on the top 100 DBLP authors and 130 principal NSF researchers with the same names; (2) DBLP–DBLP, which focuses on the 9076 DBLP authors with the same names who are known to be different people; (3) BBC–DBpedia, which focuses on 691 BBC movies and 445 DBpedia films that have similar titles; (4) DBpedia–IMDb, which focuses on 96 DBpedia movies and 101 IMDb films that have similar titles; (5) RAE–Newcastle, which focuses on 108 RAE publications and 98 Newcastle papers that are similar; and (6) Rest–Rest, which focuses on 113 and 752 restaurants published by OAEI. The scenarios DBLP–DBLP, Rest–Rest and RAE–Newcastle rely on datasets from the literature which have not being curated nor modified, the gold standard of the former two is released with the data and the rest have being extracted from LinkLion³. BBC–DBpedia, DBLP–NSF and DBPEDIA–IMDb uses a subset of the original datasets due to their size, but again data has not being modified and gold standard have being generated by hand thanks to the size reduction.

Baseline: Our baseline was the Genlink implementation from the Silk Framework 2.6.0, which is a state-of-the-art genetic programming to learn link rules.

Measures: On the one hand we explored a large portion of the parameter space to establish optimal values for θ , μ , ρ in each scenario. On the other hand we measured the number of links returned by each proposal (*Links*), precision (P), recall (R), and the F_1 score (F_1) using 2-fold cross validation. We also computed the normalised differences in precision (ΔP), recall (ΔR), and F_1 score (ΔF_1), which measure the ratio from the difference found between the baseline and our proposal and the maximum possible difference for each performance measure.

Results: The results are presented in Listing 2. We analyse them regarding precision, recall, and the F_1 score below.

The results regarding precision clearly show that our technique improves the precision of the rules learnt by GenLink in every scenario. In average, the difference in precision is 77%. The worst improvement is 57% in the DBpedia–IMDb scenario since these datasets are clearly unbalanced: movies in DBpedia have related actors, writers, or directors that IMDb may not contain at all; this obviously makes it impossible for our proposal to find enough context to make a decision. The best improvement is 100% in the DBLP–DBLP scenario since there

¹ The prototype is available at <https://github.com/AndreaCimminoArriaga/TeidePlus>.

² The datasets are available at <https://goo.gl/Pu76SU>.

³ <http://www.linklion.org/>.

are 9 069 authors with very similar names, which makes it almost impossible for GenLink to generate rules with good precision building solely on the attributes of the resources. The p-value computed by Iman-Davenport's test is 0.00; since it is clearly smaller than the standard confidence level, we can interpret it as a strong indication that there is enough evidence in our experimental data to confirm the hypothesis that our proposal works better than the baseline regarding precision.

The normalised difference of recall ΔR shows that our proposal generally retains the recall of the link rules learnt by GenLink, except in the DBpedia-IMDb and the Rae-Newcastle scenarios. The problem with the previous scenarios was that there are many incomplete resources, that is, many resources without neighbours. For instance, there are 43 papers in the Newcastle dataset that are not related to any authors. The incompleteness of data has also a negative impact on the recall of the base link rules. In our prototype, we are planning on implementing a simple check to identify incomplete resources so that the links in which they are involved are not discarded as false positives, but identified as cases on which our proposal cannot make a sensible decision. Iman-Davenport's test returns 0.25 as the corresponding p-value; since it is larger than the standard confidence level, it may be interpreted as a strong indication that the differences in recall found in our experiments are not statistically significant. In other words, the cases in which data are that incomplete do not seem to be common-enough for them to have an overall impact on our proposal.

We also studied ΔF_1 , which denotes the normalised difference in F_1 score. Note that it is 58% in average and that the corresponding Iman-davenport's p-value is 0.00, which can be interpreted as a strong indication that the difference is significant from a statistical point of view. Overall, this result confirms that our proposal helps improve precision without degrading recall and that the improvement in precision is enough for the F_1 score to improve significantly.

5 Conclusions

Data interoperability of business systems based on Web of Data requires to link the resources that are available in different datasets and represent the same real-world entities. Such links are generated by link rules that take the values of the attributes of the resources into account, but not their neighbours, which sometimes results in false positives that have a negative impact on their precision. We have presented a novel proposal that leverages a genetic programming to learn a set of link rules and then boosts them, which has proven to improve the overall F_1 score.

References

1. Alili, H., Belhajjame, K., Grigori, D., Drira, R., Ghezala, H.H.B.: On enriching user-centered data integration schemas in service lakes. In: Abramowicz, W. (ed.) BIS 2017. LNBIP, vol. 288, pp. 3–15. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-59336-4_1
2. Ananthakrishna, R., Chaudhuri, S., Ganti, V.: Eliminating fuzzy duplicates in data warehouses. In: VLDB, pp. 586–597 (2002)
3. Bhattacharya, I., Getoor, L.: Collective entity resolution in relational data. TKDD **1**(1), 1–36 (2007)
4. Bizer, C., Heath, T., Berners-Lee, T.: Linked data: principles and state of the art. In: WWW (Invited talks) (2008). <https://www.w3.org/2008/Talks/WWW2008-W3CTrack-LOD.pdf>
5. Cruz, I.F., Antonelli, F.P., Stroe, C.: AgreementMaker: efficient matching for large real-world schemas and ontologies. PVLDB **2**(2), 1586–1589 (2009)
6. Dong, X., Halevy, A.Y., Madhavan, J.: Reference reconciliation in complex information spaces. In: SIGMOD, pp. 85–96 (2005)
7. Hernández, M.A., Stolfo, S.J.: The merge/purge problem for large databases. In: SIGMOD Conference, pp. 127–138 (1995)
8. Holub, M., Proksa, O., Bieliková, M.: Detecting identical entities in the semantic web data. In: Italiano, G.F., Margaria-Steffen, T., Pokorný, J., Quisquater, J.-J., Wattenhofer, R. (eds.) SOFSEM 2015. LNCS, vol. 8939, pp. 519–530. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46078-8_43
9. Hu, W., Qu, Y.: Falcon-AO: a practical ontology matching system. J. Web Sem. **6**(3), 237–239 (2008)
10. Huber, J., Szttyler, T., Nößner, J., Meilicke, C.: CODI: combinatorial optimization for data integration. In: OM, pp. 134–141 (2011)
11. Isele, R., Bizer, C.: Learning expressive linkage rules using genetic programming. PVLDB **5**(11), 1638–1649 (2012)
12. Isele, R., Bizer, C.: Active learning of expressive linkage rules using genetic programming. J. Web Sem. **23**, 2–15 (2013)
13. Jiménez-Ruiz, E., Cuenca Grau, B.: LogMap: logic-based and scalable ontology matching. In: Aroyo, L., et al. (eds.) ISWC 2011. LNCS, vol. 7031, pp. 273–288. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-25073-6_18
14. Kalashnikov, D.V., Mehrotra, S., Chen, Z.: Exploiting relationships for domain-independent data cleaning. In: SDM, pp. 262–273 (2005)
15. Köpcke, H., Rahm, E.: Frameworks for entity matching: a comparison. Data Knowl. Eng. **69**(2), 197–210 (2010)
16. Lacoste-Julien, S., Palla, K., Davies, A., Kasneci, G., Graepel, T., Ghahramani, Z.: SIGMa: Simple greedy matching for aligning large knowledge bases. In: KDD, pp. 572–580 (2013)
17. Monge, A.E., Elkan, C.: The field matching problem: algorithms and applications. In: KDD, pp. 267–270 (1996)
18. Ngomo, A.C.N., Auer, S.: LIMES: a time-efficient approach for large-scale link discovery on the Web of data. In: IJCAI, pp. 2312–2317 (2011)
19. Ngomo, A.C.N., Lyko, K.: EAGLE: efficient active learning of link specifications using genetic programming. In: Simperl, E., Cimiano, P., Polleres, A., Corcho, O., Presutti, V. (eds.) ESWC 2012. LNCS, vol. 7295, pp. 149–163. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-30284-8_17

20. Nikolov, A., d'Aquin, M., Motta, E.: Unsupervised learning of link discovery configuration. In: Simperl, E., Cimiano, P., Polleres, A., Corcho, O., Presutti, V. (eds.) ESWC 2012. LNCS, vol. 7295, pp. 119–133. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-30284-8_15
21. Rastogi, V., Dalvi, N.N., Garofalakis, M.N.: Large-scale collective entity matching. PVLDB 4(4), 208–218 (2011)
22. Soru, T., Ngomo, A.C.N.: A comparison of supervised learning classifiers for link discovery. In: SEMANTICS, pp. 41–44 (2014)
23. Szekely, P., et al.: Building and using a knowledge graph to combat human trafficking. In: Arenas, M., et al. (eds.) ISWC 2015. LNCS, vol. 9367, pp. 205–221. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-25010-6_12
24. Volz, J., Bizer, C., Gaedke, M., Kobilarov, G.: Silk: a link discovery framework for the web of data. In: LDOW (2009)



Increasing the Explanatory Power of Investor Sentiment Analysis for Commodities in Online Media

Achim Klein^(✉), Martin Riekert, Lyubomir Kirilov, and Joerg Leukel

Information Systems 2, University of Hohenheim,
Schwerzstrasse 35, 70599 Stuttgart, Germany
achim.klein@uni-hohenheim.de

Abstract. Online media are an important source for investor sentiment on commodities. Although there is empirical evidence for a relationship between investor sentiment from news and commodity returns, the impact of classifier design on the explanatory power of sentiment for returns has received little attention. We evaluate the explanatory power of nine classifier designs and find that (1) a positive relationship holds between more opinionated online media sentiment and commodity returns, (2) weighting dictionary terms by machine learning increases explanatory power by up to 25%, and (3) the commonly used dictionary of Loughran and McDonald is detrimental for commodity sentiment analysis.

Keywords: Investor sentiment · Online media · Classifier design
Explanatory power · Commodity returns

1 Introduction

Beside fundamental news, *investor sentiment* is an important behavioral driver of asset prices such as commodities [1]. To succeed in the market, investors must include the opinions of other investors into their decision-making processes. A particular type of opinion is textual investor sentiment that is available from online media. This investor sentiment can automatically be retrieved and processed by text classifiers. The use of text classifiers for understanding the role of investor sentiment in financial markets has attracted much inquiry. With respect to commodities, present studies concern the impact of investor sentiment in a Thomson Reuters financial news dataset (TRN) on abnormal returns of crude oil and gold [2], crude oil prices [3], gold futures' returns [4], returns and volatility of natural gas [5], informed trading of crude oil [6], commodity market volatility [7], and bullish vis-a-vis bearish phases of commodity markets [8].

While the insights from prior studies are largely conclusive about the nature of the relationship between investor sentiment and commodity returns, the empirical literature lacks in understanding the impact of classifier design on the explanatory power regarding returns. Prior work either concentrates on black-box investor sentiment classification methods provided by Thompson Reuters [5] or uses dictionary-based methods only [6]. However, dictionary-based methods do not consider weighting the importance of terms.

In particular, prior research (e.g., [2, 3, 6]) often uses generic or stock-specific dictionaries (e.g., Loughran and McDonald [9]) without term weighting. This practice is in stark contrast to the important role of term weighting in automatic text retrieval [10]. Thus, explanatory power might be impaired, which then negatively affects the interpretation of findings reported in empirical studies. We address this gap in the literature by building on machine learning-based sentiment classification research [11–14].

The objective of this research is to evaluate classifier designs that complement plain dictionary-based methods used in prior research with machine learning (ML). We report on a comprehensive study using nine different classifiers and novel domain-specific corpora. The main finding is that such designs achieve up to 25% higher explanatory power regarding returns. Statistical tests show that the dictionary of Loughran and McDonald [15], used in prior research (e.g., [2]), is detrimental to high explanatory power for commodity returns. Our findings suggest that investor sentiments of other dictionaries (i.e., Henry [16]) should be used in combination with machine learning-based term weighting when studying the relationship between investor sentiments and market variables. To enable such research in the first place, we have created a novel corpus of labelled online media articles for training machine learning classifiers. Further, we provide datasets of online media articles on commodities; these datasets can be used for future empirical studies and are freely available – unlike the TRN. In contrast to the TRN, our dataset also includes blog articles, which are more subjective than news articles, and thus carry more sentiment. Furthermore, we study not only individual commodities but also the asset class of commodities.

Our paper proceeds as follows. We first discuss the theoretical background to our research. Then, we describe our methodology for the design of classifiers for investor sentiment in online media texts. We report on our evaluation of classifier designs. Finally, we discuss the implications and limitations of our study before concluding the paper.

2 Background

Investor sentiment can be defined as a bullish or bearish expectation of market participants about the future returns of an asset relative to average returns [17]. Investor sentiment can be elicited by means of surveys in a direct way [17], or indirectly by means of market data such as share turnover [18] and put-call-ratios [19]. An important source of investor sentiments is *online media* such as financial news and investment blogs on the internet. This textual source comes with the advantage of being a direct form of investor sentiment, which can be made available at high frequency and low cost. In the following paragraphs, we discuss prior research regarding (1) approaches for the textual classification of investor sentiment and (2) findings on the explanatory power of automatically classified investor sentiment from textual sources for commodity returns.

2.1 Investor Sentiment Classification Approaches

Investor sentiment classification is treated as a binary text classification task that maps a text to either a positive or a negative investor sentiment class. The positive (negative) class refers to the expectation of the author of positive (negative) future returns of a commodity or the commodities asset class discussed in their text. The aggregated sentiments of several texts form a metric, which can then be used as a sentiment index [2].

Dictionary-based approaches rely on human experts, who manually define dictionaries of words with context-free positive or negative (investor) sentiment connotation (e.g., [9, 16]). For investor sentiment classification, dictionary-based approaches extract the number of positive and negative words from a text and classify the text as positive, if the text contains more positive than negative words, and as negative otherwise. A major problem with dictionary-based approaches, which have been primarily used in information systems (IS) research (e.g., [6]), is their construct validity. Specifically, the individual importance of terms for explaining commodity returns is not considered because the weight of each term is equal. Thus, the explanatory power of sentiments based solely on dictionaries might suffer.

The state-of-the-art approach for weighting term importance for sentiment text classification is *machine learning-based* [11, 12] and should allow for higher explanatory power than plain dictionary-based approaches. The advantage of machine learning-based approaches is that they do not require collecting additional knowledge because they learn the importance of term weights automatically from training examples. Another advantage of machine learning-based methods is that they can handle high dimensional spaces of large vocabularies very well [11, 20]. Machine learning-based approaches also increase the classification accuracy by weighting features extracted from a text (e.g., [21]).

2.2 Explanatory Power for Commodity Returns

Based on arguments of behavioral finance theory, investor sentiment should have a positive effect on asset returns [17]. This effect can persist for some time because of correlated trading based on investor sentiment and limits of arbitrage that would counter investor sentiment-based trading [22–24]. Empirical evidence for these effects has been found by previous research using textual news [2, 4]. However, the impact of the design of the classifier for investor sentiment on its explanatory power and the findings regarding relationship to returns remains unclear. For better understanding of this impact, we evaluate multiple classifier designs in terms of their correlation to commodity returns. In addition, we advance prior IS research by studying a broader set of texts from online media that includes blogs and thus is more opinionated and more related to sentiment. Furthermore, previous research has found evidence for individual commodities [3–5] but not for an aggregate level of a portfolio of commodities, i.e., a commodity index. We extend this prior research by evaluating classifier impact on the relationship of investor sentiment from online media and commodity returns on the level of the commodity market as a whole - represented by a commodity index.

3 Investor Sentiment Classifier Designs

This section describes designs for classifiers of investor sentiments in online media texts using dictionary-based approaches that prevail in the extant IS literature and more advanced machine learning approaches, which allow for term weighting.

First, we designed *dictionary-based investor sentiment classifiers*: We used dictionaries of Henry [16], and Loughran and McDonald [9], separately and in combination of both. The dictionary by Loughran and McDonald [9] was developed in relation to companies' 10-K reports. The dictionary is publicly available (http://www3.nd.edu/~mcdonald/Word_Lists.html) and contains 2,329 negative words and 354 positive words. The dictionary by Henry [16] was developed in the context of earnings press releases and contains 85 negative words and 104 positive words. To combine both dictionaries, we use the union of both sets, which results in a set of 400 positive and 2,366 negative words. According to Feuerriegel and Neumann [2], the dictionary of Henry [16] is better suited for sentiment identification and studying relationships to returns than the dictionary of Loughran and McDonald [9].

Second, we designed *dictionary-based investor sentiment classifiers with preprocessing*: Extending the above approach for dictionary-based investor sentiment classification, we replicated the approach of Feuerriegel and Neumann [2] as closely as possible: (1) Stop words were removed using a list of 571 words from Lewis et al. [25], cited in Feuerriegel and Neumann ([2], p. 6). (2) Negations were taken into account in the sense of inverting the polarity of positive or negative words from the dictionary used. Following Feuerriegel and Neumann ([2], p. 6) and Dadvar et al. [26], cited in Feuerriegel and Neumann ([2], p. 6), inversions of sentiment polarity words were conducted within a sentence in the scope of three words after the word “no” and for all words in a sentence after an occurrence of one of the following words or phrases: “rather”, “hardly”, “couldn't”, “could not”, “wasn't”, “was not”, “didn't”, “did not”, “wouldn't”, “would not”, “shouldn't”, “should not”, “weren't”, “were not”, “don't”, “do not”, “doesn't”, “does not”, “haven't”, “have not”, “won't”, “will not”, “hadn't”, “had not”, and “never”. (3) The Porter stemmer was used to stem each word in all dictionaries, word lists, and texts from online media used in this research. (4) We did not use synonym merging as proposed by Feuerriegel and Neumann ([2], p. 6) because the list of synonyms merged is not disclosed in their paper.

Third, our *machine learning-based investor sentiment classifier* design follows standard approaches and configurations for textual investor sentiment classification [13, 27]. We used the whole online media texts, represented by unigrams, for classifier training on the document level. Higher level n-grams such as bigrams were not considered because the purpose of this classifier design is to evaluate machine learning-based term weighting of sentiment dictionaries containing only single words. Thus, we restrict the feature set to unigrams that are present in a dictionary of sentiment-laden words with respect to the classes positive and negative. The dictionaries used are the ones also used in the plain dictionary-based approaches above to allow for comparison. Using the dictionaries for feature selection seems reasonable because the complete feature set of unigrams is most likely not relevant for investor sentiment classification and explaining relationships to commodity returns. This assumption was corroborated by a pre-experiment in our evaluation setup that did not improve explanatory power of investor

sentiments classified by ML using all features. Thus, we used L2-normalized term occurrence feature weighting of the dictionary terms. Furthermore, we used a Support Vector Machine (SVM) for machine learning with the LIBLINEAR implementation [28] and the default configuration, i.e., L2-regularized and L2-loss dual-form SVM with linear kernel [13]. SVM was chosen for machine learning, because it achieves the highest accuracy in general text classification tasks [20] and also in bipolar sentiment classification [11, 29]. More advanced approaches like ensembles of classifiers or deep learning were not considered in this work because (1) complex approaches like deep learning allow for high model complexity, resulting in overfitting when a corpus of training examples would be used that is comparable in size to our corpus [30], and (2) the objective of this research is to evaluate the change in explanatory power of investor sentiment classified from dictionaries in combination with standard machine learning-based term weighting. To allow for this first step, before optimizing the employed ML techniques, we (had to) contribute a novel corpus specific (and publicly yet unavailable) for commodities, described in the next section.

4 Evaluation

This section reports on evaluating the explanatory power of classifier designs in terms of correlation of classified investor sentiments regarding commodity returns. Figure 1 depicts the full process as well as the corpora and datasets used: (1) Texts are collected from feeds and classified by an asset class classifier to filter the ones related to commodities, (2) a commodity classifier categorizes the remaining texts into specific commodities, (3) investor sentiment classifiers are trained on a novel investor sentiment corpus for categorizing the texts into positive or negative classes, (4) the classifiers are applied to the texts in the dataset of commodity texts and an aggregated daily score is calculated, and (5) the correlation between investor sentiment and commodity returns is determined as our main evaluation metric.

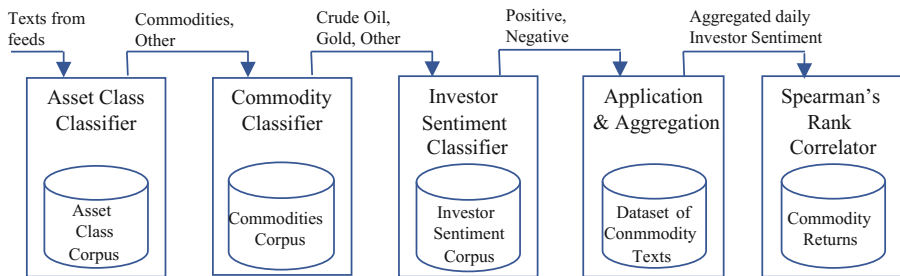


Fig. 1. The full process of training, applying, and evaluating investor sentiment classifiers.

4.1 Investor Sentiment Corpus

We created a novel investor sentiment corpus for training and evaluating investor sentiment ML classifiers. The corpus consists of texts from online media about commodities

that were manually categorized on the document level into a positive or negative investor sentiment class. To identify suitable texts, online articles related to finance and economics were manually retrieved using Google's search engine with a custom search filter such that the searched websites were the same as the websites used for the evaluation of the explanatory power of investor sentiments regarding commodity returns. For representing different market phases properly in our corpus, articles were sampled over several years. Commodity texts discussing multiple commodities (e.g., gold, crude oil) each with different investor sentiment classes were not considered for the corpus. The reason is that a uniform document level sentiment could not be assigned for the training of the document level classifiers. In our experience, such texts occur seldom. Thus, the classifiers should work reasonably well on real world datasets. Altogether, the investor sentiment corpus consists of 309 positive and 212 negative online media texts. Considering the laborious task of annotation, the corpus can be regarded as rather large.

4.2 Datasets and Preprocessing

We created a dataset of online media commodities texts to obtain investor sentiments. The dataset of commodity texts was created by crawling RSS feeds and filtering the retrieved texts for ones referring to commodities as an asset class, and the individual commodities gold and crude oil by means of an asset class classifier and a classifier for individual commodities trained on the following corpora:

The *asset class corpus* contains all texts used to train the asset class classifier. The asset class corpus is structured like a taxonomy, e.g., a text titled "Gold Should Still Shine for Investors" in the commodities asset class further belongs to the gold subclass. The corpus consists of 2,915 texts with 780 referring to commodities (of which 148 are gold, and 138 are crude oil), 1,634 referring to other asset classes (e.g., equities), and 501 *not* referring to any asset related topics, i.e., being irrelevant for this research. Texts in the commodities asset class form the *commodities corpus* and were used to train a classifier to differentiate texts that refer to different commodities.

For creating classifiers to differentiate asset classes and commodities, the default configuration for SVM was used. L2-normalized tf-idf was used for feature weighting of the unigrams, following [20, 31]. The asset class classifier achieved an accuracy of about 88% and the precision for the class of commodity texts is 93.29%, estimated in a 10-fold cross-validation on the asset class corpus. The commodity level classifier achieved an accuracy of about 98% with the following precision per class: gold (96.46%), and crude oil (98.57%). Recall is not reported because the actual commodity texts that are missed by these classifiers (as measured by recall) are not part of the subsequent actual evaluation of the explanatory power of investor sentiment classifications from specific commodity texts only.

The *dataset of commodity texts* contains 53,416 online media articles from the period October 1st, 2014 until November 13th, 2015. 22,009 texts were classified to refer to "gold", 18,580 to refer to "crude oil", and 12,827 to refer to other commodities. The articles were sampled from 450 RSS feeds of which 136 mostly refer to commodities and well represent the commodity domain. 51% of the feeds were found to contain mostly blog articles, 23% mostly news articles, and 26% about an equal amount of blog

and news articles. It must be noted, however, that the posting activity of the news feeds was in general higher than that of the blogs. The list of feeds was created by reviewing the web sites used in the EU-project FIRST [32], as well as by reviewing rankings of the top financial and economic news web sites and weblogs. Reviewed lists of websites include “50 Blogs Every Serious Trader Should Read” (<http://traderhq.com/best-trading-blogs/>) and “Top 100 Futures Trading Blogs” (<http://commodityhq.com/investor-resources/top-100-futures-trading-blogs/>).

As a first step to obtain daily investor sentiments, all commodity texts were classified by the investor sentiment classifiers on the document level. We aggregated the investor sentiment classifications to a daily score: $\ln((1 + p)/(1 + n))$ [33], where p is the number of positive text documents and n is the number of negative text documents referring to a particular commodity on a particular day (i.e., from market close to market close). On Mondays and after a bank holiday, the score incorporates also texts of the preceding weekend or bank holiday.

For relating investor sentiments to returns, *commodity log returns*, which are based on the daily closing prices of exchange-traded funds (ETFs), and a buy-and-hold strategy were used. The tickers of the ETFs used in this work are (1) DBC for the commodities asset class, (2) GLD for gold as an asset, and (3) USO for crude oil.

4.3 Correlation with Returns and Classifiers’ Accuracy

We used Spearman’s rank correlation coefficient to estimate the strength of the relationship between the daily frequency non-metric investor sentiment variable and daily buy-and-hold commodity log returns. To test for correlations, we used a standard t-test ($t_{N-2} = r_s \cdot \sqrt{N-2} / \sqrt{1-r_s^2}$), with r_s for Spearman’s rank correlation coefficient and $N = 284$ for the commodities asset class, gold, and crude oil.

Table 1 presents our results. The correlation coefficients of all classifiers are positive and statistically significantly different from zero at the 1% level for the commodities asset class, gold, and crude oil, except for (1) all dictionary-based classifiers that use only the Loughran and McDonald dictionary [9] and (2) the machine learning-based classifiers that use only the Loughran and McDonald dictionary [9] for dictionary-based feature selection.

Table 1 provides the following findings: For individual commodities, a combination of the dictionaries of Henry (H) [16] and Loughran and McDonald (LMcD) [9] with machine learning-based term weighting (ML-Dict.) observed the highest correlation with returns. For the commodities asset class, machine learning in combination with the dictionary of Henry [16] for feature selection observed the highest correlation. Thus, for high explanatory power of the classified investor sentiments for returns, the dictionaries should be complemented with machine learning-based term weighting.

Consistent with using ML for weighting the terms of dictionaries, the dictionary of LMcD used on its own for dictionary-based classification observed the lowest correlation with returns among all dictionaries. This finding is confirmed by a one-tailed Steiger’s Z-test [34]. That is, for the commodities asset class, the correlation with returns of investor sentiment classifiers using ML with the dictionary of H or a combination of H and LMcD is statistically significantly higher than using ML with the dictionary of

Table 1. Accuracy of different classifiers of investor sentiment and correlation between classified investor sentiments and commodity returns

Classifier of investor sentiment	Accuracy and micro-avg. F1 (%)	Correlation coefficients		
		Commodities	Gold	Crude oil
Dict. (LMcD)	55.1	0.075	0.139*	0.127*
Dict. P. (LMcD)	52.4	0.118*	0.221***	0.112
Dict. (H)	72.9	0.177**	0.222***	0.200***
Dict. P. (H)	71.8	0.168**	0.170**	0.187**
Dict. (H, LMcD)	69.5	0.180**	0.233***	0.194***
Dict. P. (H, LMcD)	66.6	0.173**	0.209***	0.170**
ML- Dict. (LMcD)	70.8	0.111	0.135*	0.158**
ML- Dict. (H)	74.6	0.226***	0.207***	0.205***
ML- Dict. (H, LMcD)	77.7	0.205***	0.266***	0.229***

Notes: Accuracy equals micro-averaged F1-measure (i.e., a weighted combination of both, recall and precision), thus providing an overall measure of classification performance. Correlations are Spearman's rank correlations between classified investor sentiments and daily commodity returns. Correlations are statistically significant at the *** = 0.1%, ** = 1%, * = 5% level. The length of the time series is $N = 284$. Abbreviations are: ML-Dict = machine learning with dictionary-based feature selection; Dict. = dictionary-based with dictionaries of H = Henry [16] and LMcD = Loughran and McDonald [9]; Dict. P. = dictionary-based with preprocessing [2].

LMcD alone ($p = 0.024$ and $p = 0.036$ respectively). Further, the correlation with returns is also statistically significantly higher than using a plain dictionary-based approach using LMcD on its own without preprocessing ($p = 0.017$ and $p = 0.011$ respectively) or with preprocessing ($p = 0.063$ and $p = 0.042$ respectively). The statistical evidence for individual commodities points into the same direction but is a bit weaker. Thus, the dictionary of LMcD should not be used stand-alone for high explanatory power regarding returns of commodities.

Another insight of the analysis in Table 1 is that adding pre-processing to the dictionary-based classifiers in most cases did not increase the correlation with returns. Statistical evidence, again using the one-tailed Steiger's Z-test, supports this finding for the dictionary of H and also LMcD regarding the correlation to gold returns ($p = 0.039$ and $p = 0.035$).

Consistent with above findings, the ML-based sentiment classifier observing the highest explanatory power for gold and crude oil returns also has the highest overall classification performance (in terms of accuracy and F1 measure) regarding the human-made gold standard corpus. By using ML-based classifiers, the accuracy can be improved considerably compared to plain dictionary-based approaches used in previous research. Furthermore, we did not observe that preprocessing methods [2] help to improve classifiers' accuracy.

5 Discussion

The evaluation of classifiers showed a positive correlation of investor sentiment from online media texts and returns of both, individual commodities and the commodities asset class. Thus, the classifiers might help in investors' decision-making processes. This finding relates to online media texts and thus extends prior findings on the relationship between investor sentiment in news and commodity returns (e.g., [2, 5]). Furthermore, the evaluation suggests a combination of the dictionaries of H and LMCD with machine-learning for term weighting to be most effective.

5.1 Contributions

First, research on the explanatory power of investor sentiment classifiers regarding returns of commodities is very important for empirical studies building on such classifiers. This research contributes by evaluating and making transparent the explanatory power of nine classifier designs. Second, our novel classifier design complements previously used dictionaries with machine learning for gauging the importance of dictionary terms. Thus, we improve the correlation with returns of the commodities asset class by 25% compared to prevailing plain dictionary-based approaches. Third, we corroborate previous findings of news-based investor sentiment to explain commodity returns with regard to more opinionated online media texts that comprise a large amount of blog articles with market analyses and personal investment recommendations. This finding holds not only for individual commodities but also for the asset class of commodities.

5.2 Implications for Research

Our study has several implications for future research. First, online media texts should be investigated further in terms of explanatory or predictive effects on market variables. Second, complementing dictionaries with machine-learning based term weighting should be investigated further. To this respect, the dictionary of Henry [16] and a combination of the dictionaries of Henry [16] and Loughran and McDonald [9] seems to be a good starting point. However, creating dictionaries more fitted to the vocabulary of commodity-related texts seems worthwhile because the dictionaries of Henry [16] and Loughran and McDonald [9] relate to earnings press releases and company reports, respectively. For creating a commodity-specific dictionary, corpus-based methods have already been proposed [35] and could be applied. Furthermore, pre-processing of texts like stemming, accounting for negations, and removing stop words does not systematically improve the explanatory power of investor sentiment classifiers regarding returns. To fuel research, we provide our datasets of commodity texts under <https://wi2.uni-hohenheim.de/analytics>. Thus, fellow researchers can also investigate machine learning-based methods and apply them in empirical studies.

5.3 Implications for Practice

Our findings also have implications for investors' decision making. First, investor sentiment from online media including blog articles might be of interest for investors and could complement traditionally used commercial financial news streams (e.g., Reuters or Bloomberg). Second, gold seems to be the most promising commodity studied in terms of the strength of the correlation with returns. Third, the correlation is stronger for individual commodities compared to the commodities asset class. The reason might be that the set and the weighting of commodities, represented in the proxy portfolio for the commodities asset class, is to some extent different from the commodities represented in online media feeds. Fourth, decision makers should seek investor sentiments from a classifier that incorporates machine learning.

5.4 Limitations

The limitations of our research are noted as follows. First, the time span covered by our dataset is limited to about 13 months. A larger dataset would allow for a more extensive evaluation and increase the external validity of our findings. However, this is a common limitation of related studies that use their own dataset instead of a commercial one [14, 36]. Second, the size of the corpus of human annotated texts could be larger to improve the validity of the trained investor sentiment classifiers. But again, the size of related corpora is similar (e.g., [37]) because human annotation is a costly and time-consuming task. Third, the dictionaries used in this research are not specific to commodities. Publicly available sentiment dictionaries refer to stocks or are not even specific to finance (e.g., the General Inquirer's Harvard IV dictionary). We plan to overcome this limitation with our own dictionary in a more extensive study. Fourth, we used SVM for machine learning based on previous findings that show SVM to be the best performing algorithm for sentiment text classification [13, 29]. Nevertheless, other classifiers, e.g., ensemble or deep learning classifiers, could be applied. Finally, our research does not measure actual utility of investor sentiment for selecting and timing commodities in an investor's portfolio. A portfolio simulation will be required in future work to address this issue.

6 Conclusion

Prior IS and finance research has studied investor sentiment from news using classifiers of commercial vendors or dictionary-based classifiers, which do not weigh the importance of terms regarding the problem of explaining returns. Nevertheless, weighing terms or other features is standard practice in machine learning-based classification (e.g., classification of spam). Thus, the explanatory power of classifiers that do not weigh terms might be impaired, directly affecting the results of empirical studies. To address this gap in the literature, this research evaluates nine classifier designs for better understanding the impact of classifier design on the explanatory power regarding returns. We find that complementing prevailing dictionary-based classifiers with machine learning for automatically weighing terms increases the explanatory power for the asset class of commodities by a 25% stronger correlation with returns and by a 14% stronger

correlation for individual commodities (of gold and crude oil). This result should help to improve investors' decision making and also help other IS researchers in studying relationships between investor sentiments and market variables. Furthermore, we corroborate prior findings relating to a commercial news-only dataset to hold also for more opinionated online media texts in the sense that a statistically significantly positive relationship exists between investor sentiment from online media texts and commodity returns. Thus, freely available online media texts seem to be worthwhile to be studied by investors. For fuelling further IS research in the field, we provide datasets of online media texts.

References

1. Shleifer, A.: *Inefficient Markets: An Introduction to Behavioral Finance*. Oxford University Press, New York (2000)
2. Feuerriegel, S., Neumann, D.: News or noise? How news drives commodity prices. In: *Proceedings of the 34th International Conference on Information Systems (ICIS 2013)*, Milan, Italy, pp. 1–20 (2013)
3. Feuerriegel, S., Heitzmann, S.F., Neumann, D.: Do investors read too much into news? How news sentiment causes price formation. In: *Proceedings of the 48th Hawaii International Conference on System Sciences (HICSS' 2015)*, pp. 4803–4812 (2015)
4. Smales, L.A.: News sentiment in the gold futures market. *J. Bank. Finan.* **49**, 275–286 (2014)
5. Borovkova, S., Mahakena, D.: News, volatility and jumps: the case of natural gas futures. *Quant. Finan.* **15**(7), 1217–1242 (2015)
6. Alfano, S., Feuerriegel, S., Neumann, D.: Is news sentiment more than just noise? In: *Proceedings of the 23rd European Conference on Information Systems (ECIS 2015)*, Münster, Germany, pp. 1–16 (2015)
7. Clements, A.E., Todorova, N.: Information flow, trading activity and commodity futures volatility. *J. Futures Markets* **36**(1), 88–104 (2016)
8. Feuerriegel, S., Lampe, M.W., Neumann, D.: News processing during speculative bubbles: evidence from the oil market. In: *Proceedings of the 47th Hawaii International Conference on System Sciences (HICSS 2014)*, pp. 4103–4112 (2014)
9. Loughran, T., McDonald, B.: When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. *J. Finan.* **66**(1), 67–97 (2011)
10. Salton, G., Buckley, C.: Term-weighting approaches in automatic text retrieval. *Inf. Process. Manage.* **24**(5), 513–523 (1988)
11. Pang, B., Lee, L., Vaithyanathan, S.: Thumbs up? Sentiment classification using machine learning techniques. In: *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP 2002)*, Philadelphia, PA, USA, pp. 79–86 (2002)
12. Sebastiani, F.: Machine learning in automated text categorization. *ACM Comput. Surv.* **34**(1), 1–47 (2002)
13. Wang, S., Manning, C.D.: Baselines and bigrams: simple, good sentiment and topic classification. In: *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics-Volume 2*, Jeju, South Korea, vol. 94305, pp. 90–94 (2012)
14. Rao, T., Srivastava, S.: Modeling movements in oil, gold, forex and market indices using search volume index and Twitter sentiments. In: *Proceedings of the 5th Annual ACM Web Science Conference*, Paris, France, pp. 336–345 (2013)
15. Loughran, T., McDonald, B.: Loughran and McDonald financial sentiment dictionaries (2014). http://www3.nd.edu/~mcdonald/Word_Lists.html

16. Henry, E.: Are investors influenced by how earnings press releases are written? *J. Bus. Commun.* **45**(4), 363–407 (2008)
17. Brown, G.W., Cliff, M.T.: Investor sentiment and asset valuation. *J. Bus.* **78**(2), 405–440 (2005)
18. Baker, M., Wurgler, J.: Investor sentiment and the cross-section of stock returns. *J. Finan.* **61**(4), 1645–1680 (2006)
19. Pan, J., Poteshman, A.M.: The information in option volume for future stock prices. *Rev. Finan. Stud.* **19**(3), 871–908 (2006)
20. Joachims, T.: Text categorization with support vector machines: learning with many relevant features. In: Nédellec, C., Rouveirol, C. (eds.) *ECML 1998. LNCS*, vol. 1398, pp. 137–142. Springer, Heidelberg (1998). <https://doi.org/10.1007/BFb0026683>
21. Klein, A., Altuntas, O., Riekert, M., Dinev, V.: Combined approach for extracting object-specific investor sentiment from weblogs. In: *Proceedings of the 11th International Conference on Wirtschaftsinformatik (WI 2013)*, Leipzig, Germany, pp. 691–705 (2013)
22. Shleifer, A., Summers, L.H.: The noise trader approach to finance. *J. Econ. Perspect.* **4**(2), 19–33 (1990)
23. Shleifer, A., Vishny, R.W.: The limits of arbitrage. *J. Finan.* **52**(1), 35–55 (1997)
24. Barberis, N., Thaler, R.: A survey of behavioral finance. In: *Handbook of the Economics of Finance*, vol. 1, pp. 1053–1128 (2003)
25. Lewis, D.D., Rose, T.G., Li, F., Yang, Y., Rose, T.G., Li, F.: RCV1: a new benchmark collection for text categorization research. *J. Mach. Learn. Res.* **5**, 361–397 (2004)
26. Dadvar, M., Hauff, C., de Jong, F.: Scope of negation detection in sentiment analysis. In: *Proceedings of the Dutch-Belgian Information Retrieval Workshop*, pp. 16–20 (2011)
27. Riekert, M., Leukel, J., Klein, A.: Online media sentiment: understanding machine learning-based classifiers. In: *Proceedings of the 24th European Conference on Information Systems (ECIS 2016)*, Istanbul, Turkey (2016)
28. Fan, R., Chang, K., Hsieh, C.: LIBLINEAR: a library for large linear classification. *J. Mach. Learn. Res.* **9**, 1871–1874 (2008)
29. Tang, H., Tan, S., Cheng, X.: A survey on sentiment detection of reviews. *Expert Syst. Appl.* **36**(7), 10760–10773 (2009)
30. Hastie, T., Tibshirani, R., Friedman, J.: *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, Second Edition. Springer Series in Statistics, 2nd edn. Springer, Heidelberg (2009). <https://doi.org/10.1007/978-0-387-84858-7>
31. Leopold, E., Kindermann, J.: Text categorization with support vector machines. How to represent texts in input space? *Mach. Learn.* **46**(1–3), 423–444 (2002)
32. Lombardi, P., Aprile, G., Gsell, M., Winter, A., Reinhardt, M., Queck, S.: D1.1 definition of market surveillance, risk management and retail brokerage usecases (2011)
33. Antweiler, W., Frank, M.Z.: Is all that talk just noise? The information content of Internet stock message boards. *J. Finan.* **59**(3), 1259–1294 (2004)
34. Steiger, J.H.: Tests for comparing elements of a correlation matrix. *Psychol. Bull.* **87**(2), 245–251 (1980)
35. Pröllochs, N., Feuerriegel, S., Neumann, D.: Generating domain-specific dictionaries using Bayesian learning. In: *Proceedings of the 23rd European Conference on Information Systems (ECIS 2015)*, pp. 1–14 (2015)
36. Bollen, J., Mao, H., Zeng, X.: Twitter mood predicts the stock market. *J. Comput. Sci.* **2**(1), 1–8 (2011)
37. Das, S.R., Chen, M.Y.: Yahoo! For Amazon: sentiment extraction from small talk on the Web. *Manage. Sci.* **53**(9), 1375–1388 (2007)



Comparative Analysis of the Informativeness and Encyclopedic Style of the Popular Web Information Sources

Nina Khairova¹(✉), Włodzimierz Lewoniewski², Krzysztof Węcel²,
Mamyrbayev Orken³, and Mukhsina Kuralai⁴

¹ National Technical University “Kharkiv Polytechnic Institute”, Kharkiv, Ukraine
khairova@kpi.kharkov.ua

² Poznań University of Economics and Business, Poznań, Poland
{wlodzimierz.lewoniewski,krzysztof.wecel}@ue.poznan.pl

³ Institute of Information and Computational Technologies, Almaty, Kazakhstan
morkenj@mail.ru

⁴ Al-Farabi Kazakh National University, Almaty, Kazakhstan
kuka_ai@mail.ru

Abstract. Nowadays, very often decision making relies on information that is found in the various Internet sources. Preferred are texts of the encyclopedic style, which contain mostly factual information. We propose to combine the logic-linguistic model and the universal dependency treebank to extract facts of various quality levels from texts. Based on Random Forest as a classification algorithm, we show the most significant types of facts and types of words that most affect the encyclopedic-style of the text. We evaluate our approach on four corpora based on Wikipedia, social and mass media texts. Our classifier achieves over 90% F-measure.

Keywords: Encyclopedic · Informativeness · Universal dependency
Random Forest · Facts extraction · Wikipedia · Mass media

1 Introduction

Very often the decision making depends on the information that is found in various Internet sources. Enterprises increasingly use various external big data sources in order to extract and integrate information into their own business systems [1]. Meanwhile, the Internet is flooded with meaningless blogs, computer-generated spam, and texts that convey no useful information for business purposes. Firms, organizations, and individual users can publish texts that have different purposes. The quality of information about the same subject in these texts can greatly depend on different aspects. For business purposes, however, organizations and individual users need a condensed presentation of material that identifies the subject accurately, completely and authentically. At the same time, the subject matter should be displayed in a clear and understandable manner.

In other words, the decision making should prefer texts of an encyclopedic style, which is directly related to the informativeness concept i.e. the amount of useful

information contained in a document. Obviously, the amount of knowledge that human consciousness can extract from a text has to correlate with the quality and quantity of facts in the text. Based on the definitions of an encyclopedia and encyclopedia articles [2] we can suggest that an encyclopedic text has to focus on factual information concerning the particular field, which is defined. We propose to join the use of the logic-linguistic model [3] and the universal dependency treebank [4] to extract facts of various quality levels from texts.

The model that we described in our previous studies defines the complex and simple facts via correlations between grammatical and semantic features of the words in a sentence. In order to identify these grammatical and semantic features correctly, we employ the Universal Dependencies parser, which can analyze the syntax of verb groups, subordinate clauses, and multi-word expressions in the most sufficient way.

Additionally, we take into account the employing of proper nouns, numerals, foreign words and some others provided by POS-tagging morphological and semantic types of words in the text, which can have an impact on the briefness and concreteness of particular information.

In our study, we focus on using information about the quality and quantity of facts and morphological and semantic types of the words in a text to evaluate the encyclopedic style of the text.

In order to estimate the influence of quality and quantity of factual information and semantic types of words of the text on its encyclopedic style, we decided to use four different corpora. The first one comprises Wikipedia articles which are manually divided into several classes according to their quality. The second Wikipedia corpus comprises only the best Wikipedia articles. The third corpus is The Blog Authorship Corpus¹, which contains posts of 19,320 bloggers gathered from [blogger.com](http://www.blogger.com). The corpus incorporates a total of 681,288 posts and over 140 million words - or approximately 35 posts and 7250 words per person [5]. The fourth corpus comprises news reports from various topics sections of "The New York Times" and "The Guardian", which are extracted in January 2018. We apply the Random Forests algorithm of Weka 3 Data Mining Software in order to estimate the importance of investigated features in obtained classification model.

2 Related Work

Nowadays, the problem of determining informativeness of a text has become one of the most important tasks of NLP. In recent years, many articles devoted to the solution of the problem have appeared.

Usually, informativeness of a text is considered at three levels of the linguistic system: (1) at the sentence level, (2) at the level of the discourse and (3) at the level of the entire text. In traditional linguistics, the definition of the sentence informativeness is based on the division of an utterance into two parts - the topic (or theme) and the comment (or rheme, rema). [6]. At the discourse level, the traditional approach involves the anchoring of information units or events to descriptives and interpretives within a narrative frame [7].

¹ <http://u.cs.biu.ac.il/~koppel/BlogCorpus.htm>.

Many studies determine ‘informativeness’ of text via ‘informativeness’ of words in the text or - ‘term informativeness’. Herewith, the most part of known approaches to measure term informativeness falls into the statistics-based category. Usually, they estimate informativeness of words by distributional characteristics of words in a corpus. For instance, a recent Rennie and Jaakkola’s study [8] introduced the term informativeness based on the fit of a word’s frequency to a mixture of 2 Unigram distribution.

The considerably fewer studies measure the semantics value of term informativeness. In our opinion, an interesting one is Kireyev’s research [9], which defined informativeness of term via the ratio of a term’s LSA vector length to its document frequency. More recently, Wu and Giles [10] defined a context-aware term informativeness based on the semantic relatedness between the context and the term’s featured contexts (or the top important contexts that cover most of a term’s semantics).

Most of approaches use statistical information in corpora. For instance, Shams [11] explored possibilities to determine informativeness of a text using a set of natural language attributes or features related to Stylometry—the linguistic analysis of writing style. However, he focused mainly on the search for informativeness in the specific biomedical texts.

Allen et al. [12] studied the information content of analysts report texts. Authors suggested that informativeness of a text is determined by its topics, writing style, and features of other signals in the reports that have important implications for investors. At the same time, they emphasized that more information texts are more assertive and concise. Their inferences about informativeness of a text are based on investors’ reaction to the analyst reports for up to five years.

In [13] work, informativeness analysis of language is determined using text-based electronic negotiations, i.e. negotiations conducted by text messages and numerical offers sent through electronic means. Applying Machine Learning methods allowed the authors to build the sets of the most informativeness words and n-grams, which are related to the successful negotiations.

Lex et al. guessed that informativeness of a document could be measured through factual density of a document, i.e. the number of facts contained in a document, normalized by its length [14].

Although the concept of encyclopedicness is closely related to the informativeness, it also includes such notions as brevity and correspondence to a given subject-matter. We suppose that the notion of ‘encyclopedicness of the text’ is more interesting and more useful than ‘informativeness of the text’ because it bases on knowledge concerning the particular subject-matter. Therefore, in our study, we consider the influence both of the various types of facts and semantic types of words of the text on the encyclopedic style of the text.

3 Methodology

3.1 Logical-Linguistic Model

We argue that the encyclopedic style of an article can be represented explicitly by the various linguistic means and semantic characteristics in the text. We have defined four

possible semantic types of facts in a sentence, which, in our opinion, might help to determine the encyclopedicness of the text. Figure 1 shows the structural scheme for distinguishing four types of facts from simple English sentences.

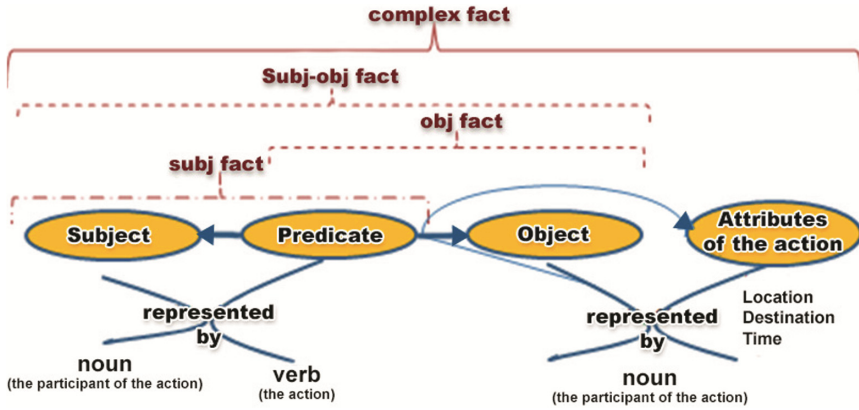


Fig. 1. Structural scheme for distinguishing four types of facts from simple English sentences: subj-fact, obj-fact, subj-obj fact and complex fact.

We called the first type of facts as *subj-fact*. We defined this type of the facts in an English sentence as the smallest grammatical clause that includes a verb and a noun. In this case, the verb represents Predicate of an action and the noun represents Subject² of an action. According to our model of fact extraction from English texts [3], the semantic relations that denote the Subject of the fact can be determined as the following logical-linguistic equation:

$$\gamma_1(z, y, x, m, p, f, n) = y^{out}((f^{can} \vee f^{may} \vee f^{must} \vee f^{should} \vee f^{could} \vee f^{need} \vee f^{might} \vee f^{would} \vee f^{out})(n^{not} \vee n^{out})(p^I \vee p^{ed} \vee p^{III})x^f m^{out} \vee (x^I(m^{is} \vee m^{are} \vee m^{havb} m^{hasb} \vee m^{hadb} \vee m^{was} \vee m^{were} \vee m^{be} \vee m^{out})z^{by}), \tag{1}$$

where the subject variable z^{prep} defines the syntactic feature of the preposition after the verb in English phrases; the subject variable y ($y^{ap} \vee y^{aps} \vee y^{out} = 1$) defines whether there is an apostrophe in the end of the word; the subject variable x defines the position of the noun with respect to the verb; the subject variable m defines whether there is a form of the verb “to be” in the phrase and the subject variable p defines the basic forms of the verb in English.

Additionally, in this study, we have appended two subject variables f and n to account for modality and negation. The subject variable f defines the possible forms of modal verbs:

$$f^{can} \vee f^{may} \vee f^{must} \vee f^{should} \vee f^{could} \vee f^{need} \vee f^{might} \vee f^{would} \vee f^{out} = 1$$

² We use ‘Predicate’, ‘Subject’ and ‘Object’ with the first upper-case letters to emphasize the semantic meaning of words in a sentence.

Using the subject variable n we can take into account the negation in a sentence:

$$n^{\text{not}} \vee n^{\text{out}} = 1$$

Definition 1. The *subj-fact* in an English sentence is the smallest grammatical clause that includes a verb and a noun (or personal pronoun) that represents the Subject of the fact according to the Eq. (1).

The Object of a verb is the second most important argument of a verb after the subject. We can define grammatical and syntactic characteristics of the Object in English text by the following logical-linguistic equation:

$$\begin{aligned} \gamma_2(z, y, x, m, p, f, n) = & y^{\text{out}}(n^{\text{not}} \vee n^{\text{out}})(f^{\text{can}} \vee f^{\text{may}} \vee f^{\text{must}} \vee f^{\text{should}} \vee f^{\text{could}} \vee f^{\text{need}} \vee \\ & f^{\text{might}} \vee f^{\text{would}} \vee f^{\text{out}})(z^{\text{out}} x^{\text{I}} m^{\text{out}} (p^{\text{I}} \vee p^{\text{ed}} \vee p^{\text{III}}) \vee x^{\text{f}}(z^{\text{out}} \vee z^{\text{by}})(m^{\text{is}} \vee m^{\text{are}} \vee \\ & m^{\text{habv}} \vee m^{\text{hasb}} \vee m^{\text{hadb}} \vee m^{\text{was}} \vee m^{\text{were}} \vee m^{\text{be}} \vee m^{\text{out}})(p^{\text{ed}} \vee p^{\text{III}}), \end{aligned} \quad (2)$$

Definition 2. The *obj-fact* in an English sentence is the smallest grammatical clause that includes a verb and a noun (or personal pronoun) representing the Object of the fact according to the conjunction of grammatical features in Eq. (2).

The third group of facts includes main clauses, in which one noun has to play the semantic role of the Subject and the other has to be the Object of the fact.

Definition 3. The *subj-obj fact* in an English sentence is the main clause that includes a verb and two nouns (or personal pronouns), one of them represents the Subject of the fact accordance with the Eq. (1) and the other represents the Object of the fact accordance with the Eq. (2).

We also defined the *complex fact* in texts as a grammatical simple English sentence that includes a verb and a few nouns (or personal pronouns). In that case, the verb also represents Predicate, but among nouns, one of these has to play the semantic role of the Subject, other has to be the Object and the others are the attributes of the action.

Definition 4. The *complex fact* in an English sentence is the simple sentence that includes a verb and a few nouns (or personal pronouns), one of them represents the Subject, another represents the Object of the fact in accordance with the Eqs. (1) and (2) respectively and the others represent attributes of the action.

These can be the attributes of time, location, mode of action, affiliation with the Subject or the Object, etc. According to our previous articles, the attributes of an action in English simple sentence can be represented by nouns that were defined by the logical-linguistic equations [12].

Additionally, we distinguish a few semantic kinds of nouns that can be extracted by labels of POS-tagging. Moreover, additionally, we distinguish a few semantic types of nouns that can be extracted by labels of POS-tagging. These are *NNP** (plural and single proper noun), *CD* (numeral), *DT* (determiner, which marked such words as “all”, “any”, “each”, “many” etc.), *FW* (foreign word), *LS* (list item marker), *MD* (modal auxiliary). The approach is based on our hypothesis that occurrence of the proper names, numerals, determiner words, foreign words, items markers, modal auxiliary words in a text can influence its encyclopedicness. For instance, the occurrence of proper nouns, numerals,

foreign words and items markers in a sentence can explicitly represent that a statement is formulated more precisely and concisely. On the contrary, we can guess that occurrence of modal auxiliary words in a sentence makes the statement vaguer and more implicit.

3.2 Using Universal Dependencies

In order to correctly pick facts out and properly distinguish their type, we employ the syntactic dependency relation. We exploit Universal Dependencies parser because for this treebanks can the most sufficiently analyze verb groups, subordinate clauses, and multi-word expressions for a lot of languages. The dependency representation of UD evolves out of Stanford Dependencies (SD), which itself follows ideas of grammatical relations-focused description that can be found in many linguistic frameworks. That is, it is centrally organized around notions of subject, object, clausal complement, noun determiner, noun modifier, etc. [4, 15]. These syntactic relations, which connect words of a sentence to each other, often express some semantic content. The verb is taken to be the structural center of clause structure in Dependency grammar and all other words are either directly or indirectly connected to it. In Dependency grammar just as a verb is considered to be the central component of a fact and all participants of the action depend on the Predicate, which expresses the fact and is represented by a verb [16]. For example, Fig. 2 shows the graphical representation of Universal Dependencies for the sentence “*The Marines reported that ten Marines and 139 insurgents died in the offensive*”, which is obtained using a special visualization tool for dependency parse - Dependencee³.




Fig. 2. Graphical representation of Universal Dependencies for the sentence “*The Marines reported that ten Marines and 139 insurgents died in the offensive*”. Source: Dependencee

For our analysis we used 7 out of 40 grammatical relations between words in English sentences, which UD v1 contains⁴. In order to pick the *subj*-fact out, we distinguish three types of dependencies: *nsubj*, *nsubjpass* and *csubj*. *Nsubj* label denotes the syntactic subject dependence on a root verb of a sentence, *csubj* label denotes the clausal syntactic subject of a clause, and *nsubjpass* label denotes the syntactic subject of a passive clause. Additionally, in order to pick the *obj*-fact out, we distinguish three types of

³ <http://chaoticity.com/dependsee-a-dependency-parse-visualisation-tool/>.

⁴ <http://universaldependencies.org/en/dep/>.

dependencies: *obj*, *iobj*, *doj* and *ccomp*. *Obj* denotes the entity acted upon or which undergoes a change of state or motion. The labels *iobj*, *doj* and *ccomp* are used for more specific notation of action object dependencies on the verb.

Considering that the root verb is the structural center of a sentence in Dependency grammar, we distinguish additional types of facts that we can extract from a text. The root grammatical relation points to the root of the sentence [15]. These are such fact types that are formed from the root verb (*root_fact*, *subj_fact_root*, *obj_fact_root*, *subj_obj_fact_root*, *complex_fact_root*) and the other ones, in which the action Predicate is not a root verb (*obj_fact_notroot*, *subj_obj_fact_notroot*, *complex_fact_notroot*).

For completeness of the study, we attribute sentences with copular verbs to a special type of facts, which we called *copular_fact*. We should do this for the following reason. Despite the fact that, as is widely known, the copular verb is not an action verb, such a verb can often be used as an existential verb, meaning “to exist”.

4 Source Data and Experimental Results

Our dataset comprises four corpora, two of which include articles from English Wikipedia. We consider texts from Wikipedia for our experiments for a few reasons. First, we assume that since Wikipedia is the biggest public universal encyclopedia, consequently Wikipedia’s articles must be well-written and must follow the encyclopedic style guidelines. Furthermore, Wikipedia articles can be divided into a different quality classes [16–18], hence the best Wikipedia’s articles have a greater degree of encyclopedicness than most other texts do. These hypotheses allow us to use the dataset of Wikipedia articles in order to evaluate the impact our proposed linguistic features on the encyclopedic style of texts.

The first Wikipedia corpus, which we called “*Wikipedia_6C*”, comprises 3000 randomly selected articles from the 6 quality classes of English Wikipedia (from the highest): Featured articles (FA), Good articles (GA), B-class, C-class, Start, Stub. We exclude A-class articles since this quality grade is usually used in conjunction with other ones (more often with FA and GA) as it was done in the previous studies [17, 18]. The second Wikipedia corpus, which is called “*WikipediaFA*”, comprises 3000 only the best Wikipedia articles that randomly are selected from the best quality class - FA.

In order to process plain texts of described above corpora, we use Wikipedia database dump from January 2018 and special framework WikiExtractor,⁵ which extracts and cleans text from a Wikipedia database dumps.

In addition, in order to compare the encyclopedic style of texts from Wikipedia and texts from other information sources, we have produced two further corpora. The first one is created on the basis of The Blog Authorship Corpus [5]. The corpus collected posts of 19,320 bloggers gathered from blogger.com one day in August 2004. The bloggers’ age is from 13 to 47 years⁶. For our purposes, we extract all texts of 3000 randomly selected bloggers (authors) from two age groups: “20s” bloggers (ages 23–27) and “30s” bloggers (ages 33–47). Each age group in our dataset has the same number of bloggers

⁵ <https://github.com/attardi/wikiextractor>.

⁶ Groups description available on the page <http://u.cs.biu.ac.il/~koppel/BlogCorpus.htm>.

(1500 each). Since bloggers have a different number of posts of various size, in our corpus we consider all texts of one blogger as a separate item. Hence we got in total 3000 items for our “*Blogs*” corpus.

The second supplementary corpus, which is called “*News*”, is created on the basis of articles and news from popular mass media portals, such as “The New York Times”⁷ and “The Guardian”⁸. For a more comprehensive analysis, we extracted an equal number of news from different topics. For our experiment, we selected 10 topics for each news source. So, we extracted 150 recent news from each topic of each news source (“The New York Times” and “The Guardian”) in January 2018. In total, we have got 3000 items for our News corpus.

Thus, we have had four corpora with the same number of items for our experiments. Table 1 shows the distributions of the analyzed texts according to the categories. By categories, we mean topics of newspaper posts in the “*News*” corpus, age groups of bloggers in the “*Blogs*” corpus and the special manual mark of assessment quality of Wikipedia articles in “*Wikipedia_6C*” corpus and “*Wikipedia_FA*” corpus.

Table 1. The distributions of the analyzed texts by our four corpora

Corpus name	Categories	Items in each category	Total number of items in corpus
Wikipedia_6C	FA, GA, B-class, C-class, Start, Stub	500	3000
Wikipedia_FA	FA	3000	3000
Blogs	“20s” blogs, “30s” blogs	1500	3000
News	Business, Health, N.Y., Opinion, Politics, Science, Sports, Tech, U.S., World topics of “The New York Times” UK news, World, Sport, Opinion, Culture, Business, Lifestyle, Technology, Environment, Travel topics of “The Guardian”	200	3000

Additionally, based on the Corpus Linguistics approaches [19], in order to compare the frequencies of linguistic features occurrence in the different corpora, we normalized their frequencies per million words. That allows to compare the frequencies of various characteristics in the corpora of different sizes.

Definition 5. The frequency of each feature in a corpus is defined as the number of the feature occurrence in the corpus divided by the number of words in –the corpus multiplied by million.

In order to assess the impact of the various types of facts in a sentence and some types of words in a text on the degree of encyclopedic text, we focus on two experiments. Both of them classify texts from Blogs, News and Wikipedia. The difference lies in the selected Wikipedia corpus. In the first experiment, we used texts from Wikipedia_6C

⁷ <https://www.nytimes.com/>.

⁸ <https://www.theguardian.com>.

corpus, which includes Wikipedia articles of different quality. We called this experiment as BNW6 model. In the other experiment we use texts from ‘Wikipedia_FA’ corpus, which only consists of the best Wikipedia articles. We called second experiment as BNWF model.

The used Random Forests classifier of the Data Mining Software Weka 3⁹ allows determining the probability that an article belongs to one of the three corpora. Table 2 shows detailed accuracy by two models respectively.

Table 2. Detailed Accuracy by models

Model	TP rate	FP rate	Precision	Recall	F-Measure	MCC	ROC area	PRC area
BNW	0.887	0.057	0.888	0.887	0.887	0.831	0.974	0.950
BNWF	0.903	0.048	0.904	0.903	0.904	0.856	0.981	0.962

Additionally, the used Data Mining Software Weka 3 allows constructing a confusion matrix that permits to visualize the performance of the model. Such matrices for both classification models are shown in Table 3. Each row in the matrix allows representing the number of instances in a predicted class, while each column represents the instances in an actual class. It makes possible to present which classes were predicted correctly by our model.

Table 3. Confusion Matrices of the obtained models.

BNW6				BNWF			
Blogs	News	Wikipedia		Blogs	News	Wikipedia_FA	
2691	274	34	Blogs	2683	283	33	Blogs
199	2575	227	News	206	2655	140	News
10	276	2714	Wikipedia_6C	24	183	2793	Wikipedia_FA

Obviously, that the best Wikipedia articles must be well-written and consequently must follow the encyclopedic style guidelines. This is confirmed by higher coefficients of recall and precision of BNWF classification model than BNW6 one.

The Random Forest classifier can show the importance of features in the models. It provides two straightforward methods for feature selection: mean decrease impurity and mean decrease accuracy. Table 4 shows the most significant features, which are based on average impurity decrease and the number of nodes using that feature.

⁹ <https://www.cs.waikato.ac.nz/ml/weka/index.html>.

Table 4. The most significant features of our models based on average impurity decrease (AID) and the number of nodes using the features (NNF)

Feature	BNW6		BNWF	
	AID	NNF	AID	NNF
root_fact	0.53	7526	0.52	6354
subj_fact_root	0.47	7614	0.48	6287
subj_fact_notroot	0.45	6537	0.45	5786
obj_fact_notroot	0.42	5678	0.41	4772
obj_fact_root	0.4	5155	0.4	5354
subj_obj_fact_root	0.38	5270	0.39	4479
complex_fact_root	0.39	3994	0.38	3882
complex_fact_notroot	0.35	5541	0.35	4413
copular_fact	0.34	3924	0.33	3254
CD	0.33	4262	0.32	3668
DT	0.32	4646	0.31	4061
FW	0.29	2083	0.31	2087
MD	0.31	4388	0.3	3702
NNP*	0.29	4893	0.29	4112
LS	0.22	775	0.23	664

5 Conclusions and Future Works

We consider the determination problem of the encyclopedic style and informativeness of text from different sources as a classification task. We have four corpora of texts. Some corpora comprise more encyclopedic texts or articles and others include less encyclopedic ones.

Our study shows that factual information has the greatest impact on encyclopedicness of text. As Fig. 1 shows, we distinguish several types of facts in the sentence. They are complex *fact*, *subj fact*, *obj fact*, *subj-obj fact* and *copular-fact*. Additionally, we highlight the main fact that is represented by a sentence.

Table 4 summarizes that the most significant features that affect the encyclopedic style of the text are (1) the frequency of the main facts (*root_fact*), (2) the frequency of the subj facts, (3) the frequency of the *obj facts* and (4) the frequency of the *subj_obj* facts in a corpus. We definite all these types of facts on the basis of our logical-linguistic model and using Universal Dependencies parser.

The Random Forest classifier, which bases on our features, allows obtaining sufficiently high recall, precision and F-measure. We provide Recall = 0.887 and Precision = 0.888 in the case of the classification of texts by Blogs, News and Wikipedia corpora. In the case of considering only the best of Wikipedia articles in the last corpus, we provide recall = 0.903 and precision = 0.904. Moreover, using the Random Forest classifier allowed us to show the most important features related to informativeness and the encyclopedic style in our classification models.

In future work, we plan to extend obtained approach to compare the encyclopedic style of texts of Wikipedia and of various Web information sources in different languages. In our opinion, it is possible to implement the method in commercial or corporate search engines to provide users with more informative and encyclopedic texts. Such tools must be significant for making important decisions based on the text information from the various Internet sources. On the other hand, firms and organizations will get the opportunity to evaluate the informativeness of the texts that are placed on their Web sites, and make changes to provide more valuable information to potential users. Additionally, more encyclopedic texts can be used to enrich different open knowledge bases (such as Wikipedia, DBpedia) and business information systems in enterprises.

References

1. Cai, L., Zhu, Y.: The challenges of data quality and data quality assessment in the big data era. *Data Sci. J.* **14**, 1–10 (2015)
2. Béjoint, H.: *Modern Lexicography: An Introduction*, pp. 30–31. Oxford University Press (2000)
3. Khairova, N., Petrasova, S., Gautam, A.: The logical-linguistic model of fact extraction from English texts. In: *International Conference on Information and Software Technologies, Communications in Computer and Information Science, CCIS 2016*, pp. 625–635 (2016)
4. Nivre, J., et al.: *Universal dependencies v1: a multilingual treebank collection* In: *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016)*, Paris, France, May. European Language Resources Association (ELRA) (2016)
5. Schler, J., Koppel, M., Argamon, S., Pennebaker, J.: Effects of age and gender on blogging. In: *Proceedings of 2006 AAAI Spring Symposium on Computational Approaches for Analyzing Weblogs*, pp. 191–197 (2006)
6. Leafgren, J.: *Degrees of explicitness: information structure and the packaging of Bulgarian subjects and objects*. John Benjamins, Amsterdam & Philadelphia (2002)
7. Berman, R.A., Ravid, D.: Analyzing narrative informativeness in speech and writing. In: Tyler, A., Kim, Y., Takada, M. (eds.) *Language in the Context of Use: Cognitive Approaches to Language and Language Learning*. Cognitive Linguistics Research Series. pp. 79–101. Mouton de Gruyter, The Hague (2008)
8. Rennie, J.D.M., Jaakkola, T.: Using term informativeness for named entity detection. In: *Proceedings of SIGIR 2005*, pp. 353–360 (2005)
9. Kireyev, K.: Semantic-based estimation of term informativeness. In: *Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the ACL*, pp. 530–538 (2009)
10. Wu, Z., Giles, L.C.: Measuring term informativeness in context. In: *Proceedings of NAACL 2013*, Atlanta, Georgia, pp. 259–269 (2013)
11. Shams, R.: *Identification of informativeness in text using natural language stylometry*. Electronic Thesis and Dissertation Repository, 2365 (2014)
12. Huang, A.H., Zang, A.Y., Zheng, R.: Evidence on the information content of text in analyst reports. *Acc. Rev.* **89**(6), 2151–2180 (2014)
13. Sokolova, M., Lapalme, G.: How much do we say? Using informativeness of negotiation text records for early prediction of negotiation outcomes. *Group Decis. Negot.* **21**(3), 363–379 (2012)

14. Lex, E., Voelske, M., Errecalde, M., Ferretti, E., Cagnina, L., Horn, C., Granitzer, M.: Measuring the quality of web content using factual information. In: Proceedings of the 2nd joint WICOW/AIRWeb Workshop on Web Quality, pp. 7–10. ACM (2012)
15. De Marneffe, M.C., Manning, C.D.: Stanford typed dependencies manual, pp. 338–345. Technical report. Stanford University (2008)
16. Lewoniewski, W.: Enrichment of information in multilingual wikipedia based on quality analysis. In: Abramowicz, W. (ed.) BIS 2017. LNBIP, vol. 303, pp. 216–227. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-69023-0_19
17. Węcel, K., Lewoniewski, W.: Modelling the quality of attributes in wikipedia infoboxes. In: Abramowicz, W. (ed.) BIS 2015. LNBIP, vol. 228, pp. 308–320. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-26762-3_27
18. Lewoniewski, W., Węcel, K., Abramowicz, W.: Quality and importance of wikipedia articles in different languages. In: Dregvaite, G., Damasevicius, R. (eds.) ICIST 2016. CCIS, vol. 639, pp. 613–624. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46254-7_50
19. McEnery, T., Hardie, A.: Corpus Linguistics: Method, Theory and Practice, pp. 48–52. Cambridge University Press, Cambridge (2012)

Applications, Evaluations, and Experiences



Satellite Imagery Analysis for Operational Damage Assessment in Emergency Situations

German Novikov, Alexey Trekin, Georgy Potapov^(✉), Vladimir Ignatiev,
and Evgeny Burnaev

Skolkovo Institute of Science and Technology, Nobel st. 1, Moscow, Russia
aeronetlab@skoltech.ru, georgy.potapov@gmail.com
<http://crei.skoltech.ru/cdise/aeronet-lab/>

Abstract. When major disaster occurs the questions are raised how to estimate the damage in time to support the decision making process and relief efforts by local authorities or humanitarian teams. In this paper we consider the use of Machine Learning and Computer Vision on remote sensing imagery to improve time efficiency of assessment of damaged buildings in disaster affected area. We propose a general workflow that can be useful in various disaster management applications, and demonstrate the use of the proposed workflow for the assessment of the damage caused by the wildfires in California in 2017.

Keywords: Remote sensing · Damage assessment · Satellite imagery
Deep learning · Emergency response · Emergency mapping

1 Introduction

In emergency situations like devastating wildfires, floods, earthquakes or tsunamis, decision makers need to get information about possible damages in residential area and infrastructure most rapidly after the event. One of the most valuable sources to get such information from, is the Earth Observation systems, which include satellite and aerial remote sensing, since it can be captured shortly after the disaster without all the risks related to the ground observations. The combination of this information with statistical and ground observation data contributes to even a better valuation of physical and human losses caused by disaster [2].

There are several international programs that are arranged to support the information exchange during and after disasters such as UNOSAT (UNITAR Operational Satellite Applications Program) [4], Space disaster charter [1], Humanitarian Openstreetmap Team (HOT) [3] or Tomnod [5] which is the Digitalglobe satellite company crowdsourcing platform. All these are useful initiatives providing tools and activation workflows for emergency mapping done by specialists or by volunteers in a collaborative way [8, 9].

This method of mapping of the imagery could be time-consuming since it requires some qualification and takes time to digitize all the damages manually, particularly if the affected area is quite large and the objects are relatively small and scattered as it is for private houses in the residential area. For example, Californian wildfires past year caused significant damages in Ventura and Santa Barbara counties. The fires destroyed at least 1,063 and 5,643 structures in these areas respectively [14]. The significant delay in time of emergency mapping also might be caused by the availability of remote sensing data which has its physical limitations (cloudiness, day time, resolution etc.) as well as commercial ones (terms of use, costs etc.) Needless to say that in the post-disaster recovery strategy the time is the key factor. That's why we consider apply the Machine Learning and Computer Vision approach to the processing of Satellite and Aerial imagery to detect main damages and reduce the time costs.

2 Existing Approaches and Solutions

Among the existing solutions for emergency mapping of disaster areas it's worth to mention HOT that allows mapping of the chosen area in a collaborative way. Taking data for different areas, where HOT campaigns were activated, we estimated that the mapping process even in the areas of emergency that attract a lot of public attention, like Nepal earthquakes, takes several months to achieve the whole coverage (see Fig. 1).




Fig. 1. Time distribution of the building features created by OSM users for Kathmandu region (source - osm-analytics.org)

Following the news of this incident we found several related media publications that provide assessments of damages. We assume that the work was doing manually on satellite images - comparing to the date of incident (Dec. 4) it's taken about six days to prepare maps for the article (Dec. 11 LA Times article update) [15]. That most probably is caused by the amount of work needed to find appropriate data sources and make a damage map.

UNOSAT Rapid mapping service [4] is a framework of United Nations Institute for Training and Research in the field of emergency mapping. Even though it's quite challenging to estimate the time efficiency by their results since there is no clear definition of what states for "rapid". Usually the temporary lag of UNOSAT maps is two days from the satellite imagery acquisition date.

One of the keys to the solution of the problem might be a deep learning approach. In the last few years the deep convolutional networks became a robust solution for many problems concerning image analysis. Different variants of the networks are able to solve the problems of semantic segmentation, object classification and detection [17,20–22].

The main drawback of this class of methods is that the deep convolutional networks need a big amount of training (previously manually marked ground truth) data. On the one hand, we can pre-train the method using the data about the other event of the kind, that took place in the past. But the results of this kind of training may be unpredictable due to the difference between the data in the training and test cases. These issues are concerned in our workflow that is proposed in the following section.

In previous research the problem of organization of the workflow for rapid response to disasters using space imagery is dealt with. For example, in voigt the process of disaster mapping is considered on many different cases (fires, earthquakes etc.). However, the common issue in the techniques proposed in the paper, as well in jayce, is that they focus on the remote sensing data itself and rely on manual computer-aided mapping and procedures of image processing. These workflows cannot apply the deep learning techniques because of the training issues that have to be solved in the process. Procedure proposed in barnes where the supervisor iteratively corrects the image processing results until the machine learning method is trained. We propose a similar approach that is constructed with the deep learning in mind and thus allows us to use the most advanced methods for automatic mapping of the disaster areas.

3 Problem Statement and Proposed Workflow

The main problem we want to deal with is to decrease the time needed to retrieve crucial information for decision making in emergency situations when the proper remote sensing data is available. We propose the following workflow:

1. Determine the case of interest. The deep learning methods work significantly better when the objects of interest are similar to each other, so the case should be narrow, for example "burned houses" or "flooded roads".
2. Create a training dataset. The deep learning methods need a training data so they could learn the objects of interest and their properties. The training dataset consist of the real data (in our case, two aerospace images, one taken before the catastrophic event, and the other - after the event) and the labels that annotate and outline every damaged structure of the type.

3. Train and validate a deep learning method using the dataset. The method (or model) extracts the information from the training dataset. Its ability to detect the damaged objects of interest is validated using a subset of the training data. This pre-trained model will be used in every case of the forthcoming emergency of the given type.
4. Obtain information of a new emergency case. This is where our method starts working. The data should be of the same or similar type (spatial resolution, target objects, color depth) as that used for training, this is a critical requirement for the model to work properly.
5. Fine-tune the model for the new case. Despite the similarity of the data, the model may be unable to process them correctly due to small differences, for example different sunlight. The fine-tuning can be done using automatically annotated data from the new case, or using the manual markup for a small subset of the data.
6. Run the automatic image processing. Now that the model is trained for the case, we can make an automated processing of the rest of the data and have them ready for the decision making.

Using this approach, we need to spend some time for creation of the reference training dataset, but normally it should be made before the emergency event. Then, after the event, the amount of work needed is much less that allows us to propose a fast working and thus efficient solution.

4 The Experiment

To validate and demonstrate the workflow, we have chosen the case of wildfires in two areas of Ventura and Santa Rosa counties, California, USA, where many houses were burned to the ground in 2017. The advantage of this choice is justified by the availability of hi-resolution data provided by Digitalglobe within their Open Data Program [13]. In the following section we will follow our workflow on the case, and describe both general principles of the deep learning application to the imagery and our steps in the particular case.

4.1 Determine the Case of Interest

In our research the case of interest is “houses destroyed by fire”. A typical example of the object of interest is depicted in Fig. 2.

It is worth noting, that the case should be restricted as narrow as it is possible for it makes a big difference when speaking about the deep learning methods. For example, if we train the method on the images like this, where the houses are completely destroyed, it will not be able to detect partially damaged buildings. Also the type of building and even the rooftop material can change the result significantly.

4.2 Create a Training Dataset

The training area is chosen in the Ventura, Santa Barbara, California, that was severely affected by the Thomas Fire in the December, 2017 (see Fig. 3).

The preferable source for high-resolution data is the Digitalglobe Open Data Program. This program is one of the few sources of free high-resolution data, and the data is distributed early after the event [13]. However, in the case of South California the existing Openstreetmap (OSM) mapping which was used as the initial input for the markup is based on Google Maps/Bing Maps imagery that is better orthorectified, so that the image coordinates differ, as it can be seen in the Fig. 4. This makes existing map not as good source of the ground truth.



Fig. 2. Satellite images of buildings before and after the fire event

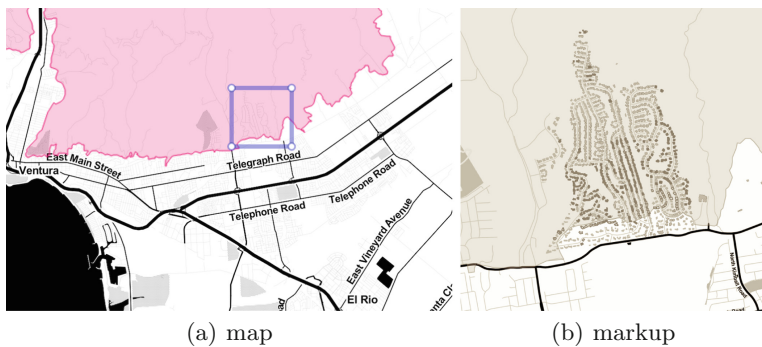


Fig. 3. Training area in Ventura and the resulting markup (Openstreetmap, Stamen design; Digitalglobe. Bounding box coordinates: -119.2278 34.2836 -119.1916 34.3065)



Fig. 4. Misalignment of the Digitalglobe image with OSM markup

Due to these reasons, we had to use the Google Maps imagery, that is similar to the Digitalglobe open data in terms of image characteristics (similar spatial resolution, also 3-band 8-bit per band images), but both pre-event and post-event images are available, and the better orthorectification leads to good alignment with the OSM.

The crowdsourced mapping data from OSM were not full and did not contain the information about burned buildings, so it was necessary to improve the OSM markup before using it as the ground truth for training the deep convolutional network. We facilitated the manual work by using of OSM as the source of initial vector features, selecting all the ones tagged as “building”. All the extracted features than were checked through the visual inspection and annotated with the appropriate additional tag “disaster” = “damaged_area” if the one was destroyed by the fire. To complete training dataset we used cartographic tools as Openstreetmap ID which is open source editor for online mapping for OSM [6]. The final markup contains 760 not damaged buildings and 320 ruined buildings (see Fig. 3) was exported in GeoJSON format using OSM API and additionally processed using our Python stack tool to convert and rasterize vector data into 1-band pixel masks.

4.3 Train and Validate a Deep Learning Method Using the Dataset

We used a semantic segmentation approach to the change detection. The semantic segmentation of an image results in a mask of the pixels that are considered to be of the target class or classes. In our case, when we have two images - before and after the event - we can gain maximum from the given data if we stack them together and make a 6-band image (3 bands before and 3 bands after). A convolutional network for change detection was built in the encoder-decoder manner, which has great success in solving semantic segmentation problems [17]. For a model that works with pairs of 3-band images, one could use a single 6-channel encoder, but this would not allow the use of a transfer-learning technique to speed up learning and improve the final quality of the results, so the model was built on a two-stream encoder, each of which “looked at its own” 3-band image and one common decoder. This approach made it possible to use the pre-trained on “ImageNet” classification dataset [16] weights for the RGB images independently in each of the branches of the encoder.

Validation on the part of the Ventura dataset that was not used for training gave appropriate results, see Fig. 5. Pixel-wise $F1$ -score for the class of burned buildings is 0.859 and for the class of unburned buildings is 0.818.




Fig. 5. Results of the change detection on the validation subset of data in Ventura. Left: image taken before fires, center: image taken after fire, right: segmentation results black - non-damaged, gray - damaged buildings

Overall training of the model has taken less than an hour on a Tesla P100 GPU using Keras with Tensorflow backend [16,18]. The trained model is now ready for the new cases of the massive fire.

4.4 Obtain Information of a New Emergency Case

We consider the fire in Santa Rosa, California (Tubbs Fire, October 2017) as the “new case” of the same type of events (see Fig. 6). The Open Data program has images both before and after the fire event, so we can use them for the test.

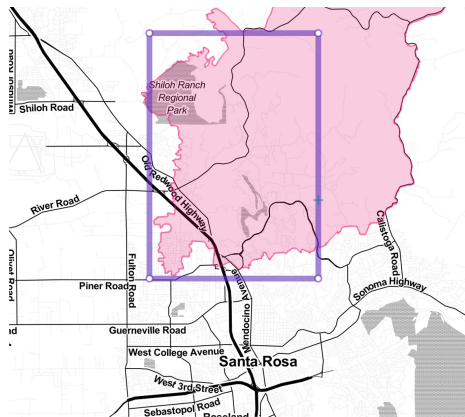


Fig. 6. A map of the new zone of the Tubbs fire in Santa Rosa, California

As the data in this case have similar characteristics, we tried the image segmentation with the model as is, without any changes. The result is unsatisfactory, however it does have some correlation with the real data. This can be caused by differences in season, solar angle, image preprocessing difference, or by some difference in the structure of the residential areas themselves. For example, buildings in Santa Rosa are closer to each other. The example of the results see in Fig. 7.




Fig. 7. Results of the change detection on the test subset of data in Santa Rosa without fine-tuning. Left: image taken before fires, center: image taken after fire, right: segmentation results black - non-damaged, gray - damaged buildings

4.5 Fine-Tune the Model for the New Case

The results above show that we need to train the model for the new area. In order to do this, we make a new small dataset in a part of the Santa Rosa, see Fig. 8. It contains 146 burned and 137 undamaged houses, so it requires far less



Fig. 8. Small dataset for fine-tuning of the net to the new data. White - unburned buildings, black contours - burned buildings

time and effort. The preparation of the dataset took about an hour of manual markup by one person.

Switching from one part of dataset to another, the results of the model were greatly deteriorated, but the dense marking of just less than 300 houses on new images allowed to improve the quality on the whole new data and reach almost the same result for 10 min of additional training.

4.6 Run the Automatic Image Processing

The rest of the Santa Rosa region of interest was processed automatically by the trained model. The example of the result taken from the test zone in the center of Santa Rosa town is shown in Figs. 9 and 10. It can be clearly seen that non fine-tuned method tends to merge the regions of the separate buildings into one area, while after the fine-tuning the resulting regions can be easily separated at the post-processing stage.




Fig. 9. An example of the test area image before and after the fire

After the fine-tuning, the change detection method can give very good results on image segmentation, and even give a good chance to distinguish between separate houses that is very important in the task of damage assessment when it is necessary to estimate the accurate number of damaged structures.

Note that the segmentation approach is more robust than the detection one because it allows to estimate the area of the target changes, that can be necessary in other emergency cases like floods, blockages, fire damage to crops or forests etc.

5 Time Efficiency

The manual markup of our Ventura training area (Fig. 3) should take about 1.5–2 days by a qualified specialist, assuming that mapping of buildings features

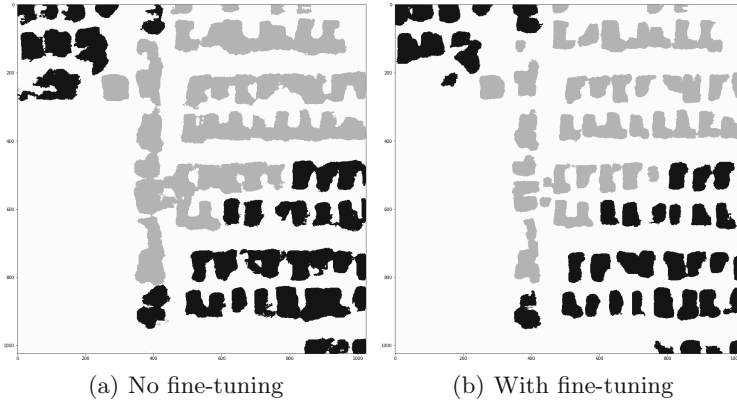


Fig. 10. Results of the image segmentation before and after fine-tuning

takes averagely 30 s per feature. But more realistic is the time evaluation of HOT mapping as it represents the real rate of community. Besides, the HOT tools are built on the top of OSM Data Base and are planned not for mapping of the state of the objects like “burnt buildings” but to improve the basic maps when the cartographic data is missed and needed by humanitarian and emergency organizations.

The full workflow for the new area, where we had to make the only a small training subset took about 3 h including model training and automatic processing.

That gives us less time needed for information retrieval for the emergency management.

6 Further Research

At the current stage we have developed a workflow and a method of the damaged areas segmentation. In the further research we will continue the development of the segmentation method to increase its accuracy and robustness to the data characteristics changes.

The method can be also extended to the problem of instance segmentation, that is distinguishing between separate objects, counting the objects, and converting them to the map that can be used online and in GIS applications.

We will apply the approach to the Open Data in the case of new events of this domain, the other types of disasters such as floods and tsunami, and will extend the training dataset to extrapolate this approach to the other cases and territories.

7 Conclusion

We have formulated the problem based on the research of the tools and frameworks for disaster mapping. Based on the problem, we proposed a workflow involving deep learning and use of open data for the emergency mapping.

We have created the training and test datasets for California fires, which means the raster mask of the vector features of damaged and non damaged buildings in the area and the appropriate pre- and post-event imagery to develop a change detection method and validate the approach.

We developed a method of change detection, based on convolutional neural networks, that is able to make semantic segmentation of the area subjected to massive fires, mapping burned buildings. The method can be easily extended to the new areas and data sources with a little training for the data peculiarities (fine tuning). We are going to elaborate on this approach by increasing its computational capabilities with large scale sparse convolutional neural networks [24], using a new loss function, specially tailored for change detection in sequences of events [25, 26], imposing a confidence measure on top of the change detector based on non-parametric conformal approach [27, 28].

The workflow turned up to give substantial profit in terms of time needed for emergency mapping and in the future we will extend it to the other cases.

Acknowledgements. The research of Evgeny Burnaev was partially supported by the Russian Foundation for Basic Research grants 16-01-00576 A and 16-29-09649 of m.

References

1. About the Disasters Charter. International Disasters Charter. 2000–2018. <https://disasterscharter.org/web/guest/about-the-charter>
2. Eguchi, R.T., et al.: The January 12, 2010 Haiti earthquake: a comprehensive damage assessment using very high resolution areal imagery. In: 8th International Workshop on Remote Sensing for Disaster Management. Tokyo Institute of Technology, Japan (2010)
3. About – Humanitarian OpenStreetMap Team. <https://www.hotosm.org/about>
4. Maps and Data – UNITAR. <http://www.unitar.org/unosat/maps>
5. Tomnod (2018). <https://www.tomnod.com/>
6. OpenStreetMap/iD: The easy-to-use OpenStreetMap editor in JavaScript (2018). <https://github.com/openstreetmap/iD>
7. Meier, P.: Human computation for disaster response. In: Michelucci, P. (ed.) Handbook of Human Computation, pp. 95–104. Springer, New York (2013). https://doi.org/10.1007/978-1-4614-8806-4_11
8. Lang, S., et al.: Earth Observation for Humanitarian Assistance GI Forum 2017, vol. 1, pp. 157–165 (2017)
9. Dittus M., Quattrone G., Capra L.: Mass participation during emergency response: event-centric crowdsourcing in humanitarian mapping. In: CSCW, 2017, pp. 1290–1303 (2017)
10. Voigt, S., Kemper, T., Riedlinger, T., Kiefl, R., Scholte, K., Mehl, H.: Satellite image analysis for disaster and crisis-management support. *IEEE Trans. Geosci. Remote Sens.* **45**(6), 1520–1528 (2007)
11. Joyce, K.E., Belliss, S.E., et al.: A review of the status of satellite remote sensing and image processing techniques for mapping natural hazards and disasters. *Progress Phys. Geogr.* **33**(2), 183–207 (2009)
12. Barnes, C.F., Fritz, H., Yoo, J.: Hurricane disaster assessments with image-driven data mining in high-resolution satellite imagery. *IEEE Trans. Geosci. Remote Sens.* **45**(6), 1631–1640 (2007)

13. Open Data Program - DigitalGlobe (2018). <https://www.digitalglobe.com/opendata>
14. http://www.fire.ca.gov/communications/downloads/fact_sheets/Top20_Destruction.pdf
15. Before and after: Where the Thomas fire destroyed buildings in Ventura (analysis made by The Los Angeles Times using sat. imagery by Digitalglobe). <http://www.latimes.com/projects/la-me-social-fires-destroyed-structures/>
16. Chollet, F.: Keras (2015)
17. Ronneberger, O., et al.: U-Net: convolutional networks for biomedical image segmentation. <https://arxiv.org/abs/1505.04597>
18. Abadi, M., Barham, P., Chen, J., et al.: TensorFlow: a system for large-scale machine learning. In: OSDI, vol. 16, pp. 265–283
19. Deng, J., Dong, W., Socher, R., Li, L. J., Li, K., Fei-Fei, L.: ImageNet: a large-scale hierarchical image database. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2009, pp. 248–255 (2009)
20. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 3431–3440 (2015)
21. Krizhevsky, A., Sutskever, I., Hinton, G.E.: ImageNet classification with deep convolutional neural networks. In: Advances in Neural Information Processing Systems, pp. 1097–1105 (2012)
22. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask R-CNN. In: 2017 IEEE International Conference on Computer Vision (ICCV), pp. 2980–2988. IEEE, October 2017
23. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint [arXiv:1409.1556](https://arxiv.org/abs/1409.1556) (2014)
24. Notchenko, A., Kapushev, Y., Burnaev, E.: Large-scale shape retrieval with sparse 3D convolutional neural networks. In: van der Aalst, W.M.P., et al. (eds.) AIST 2017. LNCS, vol. 10716, pp. 245–254. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-73013-4_23
25. Burnaev, E., Koptelov, I., Novikov, G., Khanipov, T.: Automatic construction of a recurrent neural network based classifier for vehicle passage detection. In: Proceedings of SPIE, 9th International Conference on Machine Vision (ICMV 2016), vol. 10341, p. 1034103 (2016)
26. Artemov, A., Burnaev, E.: Ensembles of detectors for online detection of transient changes. In: Proceedings of SPIE, Eighth International Conference on Machine Vision, vol. 9875, p. 98751Z:5 (2015)
27. Volkhonsky, D., Burnaev, E., Nouretdinov, I., Gammerman, A., Vovk, V.: Inductive conformal martingales for change-point detection. In: The 6th Symposium on Conformal and Probabilistic Prediction with Applications (COPA) 2017, Proceedings of Machine Learning Research, vol. 60, pp. 132–153 (2017)
28. Safin, A., Burnaev, E.: Conformal kernel expected similarity for anomaly detection in time-series data. Adv. Syst. Sci. Appl. **3**, 23–34 (2017)



Qualitative Assessment of Machine Learning Techniques in the Context of Fault Diagnostics

Thilo Habrich, Carolin Wagner^(✉), and Bernd Hellingrath

Department of Information Systems, Westfälische Wilhelms-Universität Münster,
48149 Münster, Germany

T-Habrich@web.de, {wagner,hellingrath}@ercis.uni-muenster.de

Abstract. Nowadays, in the light of high data availability and computational power, Machine Learning (ML) techniques are widely applied to the area of fault diagnostics in the context of Condition-based Maintenance (CBM). Those techniques are able to learn intelligently from data to build suitable classification models, which enable the labeling of unknown data based on observed patterns. Even though plenty of research papers deal with this topic, the question remains open, which technique should be chosen for a specific problem. In order to select appropriate methods for a given problem, the problem characteristics have to be assessed against the strengths and weaknesses of relevant ML techniques. This paper presents a qualitative assessment of well-known ML techniques based on criteria obtained from literature. It is completed by a case study to identify the most suitable techniques to perform fault diagnostics in in-vitro diagnostic instruments with regard to the presented qualitative assessment.

Keywords: Machine learning · Condition-based maintenance · Fault diagnostics
Condition monitoring

1 Introduction and Current Research Gaps

The preventive CBM strategy enables planning and execution of maintenance actions based on the actual maintenance demand of machines. It is based on the principle of continuously monitoring the machine condition using sensors. Two major tasks are performed by CBM, namely the identification of fault types (fault diagnostics) and the estimation of the remaining useful life (prognostics).

The selection of the most suitable ML techniques for specific fault diagnostics tasks is often elusive and based on gut feeling. The performance of each technique is highly dependent on the specific problem, in particular considering data characteristics as well as the external inputs. Individual ML techniques, however, exhibit several general characteristics, which make them appropriate to a greater or lesser extent for a specific problem. By assessing and weighting these specific characteristics for a certain problem, the selection of a suitable algorithm can be supported. This paper aims to provide a qualitative assessment of ML techniques for fault diagnostics, which facilitates the selection of a suitable technique for a specific problem at hand. Based on a structural literature research, ML techniques are identified which have been applied to fault

diagnostics in CBM. In a subsequent step, those ML techniques are qualitatively analyzed, adhering to several general requirements for ML techniques. The assessment is based on existing studies and extended with the identified algorithms and requirements for fault diagnostics.

The paper is structured as follows: Sect. 2 provides an overview of CBM and ML, especially highlighting machine-learning techniques applied to condition-based maintenance. In Sect. 3, a requirement catalogue for the assessment of machine learning techniques in condition-based maintenance is developed. This is followed in Sect. 4 by an industry case applying and validating the developed requirement catalogue.

2 State-of-the-Art

In order to present a short overview of the covered concepts first CBM and ML are introduced. The last section shows how ML has been applied to CBM in the past.

2.1 Condition-Based Maintenance

Evolving from simpler maintenance techniques like corrective and preventive maintenance, CBM has been the focus of research for a larger period of time [1]. It was introduced to maximize the uptime of machines by avoiding unplanned breakdowns in combination with lowering maintenance costs. CBM tries to detect, predict and respond to machinery failures before they occur, by assessing the machinery's condition, based on operational data [2]. CBM is initiated by the fact that most faults are preceded by certain symptoms, indicating a degradation process. However, CBM analysis has to cope with several difficulties. These difficulties include dealing with a large number of continuously observed process variables (sometimes even more than 1500 process variables). In addition, difficulties are imposed by insufficient, unreliable or incomplete data introduced by sensor malfunctions or incomplete labeling, which is sometimes performed by hand [3]. Therefore, systems need to be able to deal with missing values.

To come to maintenance decisions, based on the degradation process, CBM consists of two major elements, diagnostics and prognostics. The diagnostics process is used to detect and diagnose faults [4], while prognostics predicts faults and deterioration of components [5]. In the following, the paper will focus on diagnostics.

Diagnostics. Diagnostics, whose final process goal is accurate fault detection and diagnosis (FDD), can be viewed as the process of obtaining the relationship between cause and effect of a fault [6]. Hence, it includes the tasks: fault detection, isolation and identification. Fault detection continuously analyzes whether and when a fault occurred [7]. The faulty component is then identified during fault isolation [8]. According to Schwabacher and Goebel [9], fault identification “requires determining the specific fault mode, rather than just reporting which sensor has an unusual value” (p. 107). The fault mode is an estimate of nature and extent of a fault [8]. Determining fault modes is important in order to take failure-tolerant actions and thereby suitably restore machines' normal condition [7]. Ultimately, diagnostics tries to solve classification tasks, because

it attempts to separate data into normal and faulty states, including the type and location of faults [10].

During machine operation, several processes are running on a machine. Condition monitoring (CM) monitors outputs from sensors created during the machine's processes out of known inputs [11]. Methods for fault diagnostics can be grouped into three classes [7].

The first class is called model-based, because its methods use a mathematical model of the physical system. Methods in this class obtain their models by using physical principles or system identification techniques [7]. Using these methods, only small amounts of online data - recorded during CM - are needed to be compared to the expected output generated by the model, using the same inputs as the machine's process [11]. The second class, signal-based methods, analyzes continuous signals, measured from sensors by looking for anomalies without using an input-output model [7]. These methods can be applied to either online or historic data [11].

Diagnostics has to be performed using learn-by-example methods, in cases where the analyzed system is too complex to allow model creation or is not displaying anomalies within one continuous measured signal. These methods are part of the third class of knowledge-based methods. Recent knowledge-based diagnostics techniques analyze whether consistency exists between current process data and implicit knowledge obtained from huge amounts of multisource historical process data [11]. Due to recent research advances in this field, the focus of this paper is set on this area of methods.

2.2 Machine Learning in the Context of Artificial Intelligence

Generally, ML is considered a subsection of Artificial Intelligence (AI), even though its scope differs slightly from the scope of AI. According to Marsland [13], ML "is about making computers modify or adapt their actions [...] so that these actions get more accurate" (p. 4). AI on the other hand describes, according to the ISO/IEC-STANDARD 2382:2015 (No. 2121393) [14], a machines' ability to conduct a set of functions such as reasoning and learning that are usually linked to human intelligence. ML is becoming very popular, because it avoids the necessity to build an explicit (mathematical) model of the analyzed subject. In FDD, ML is used to analyze process data, which was recorded during machine operation to extract process-specific knowledge [11]. Effectiveness of ML derives from its ability to perform generalization [13]. This effectiveness comes along with the costs of extensive computational power. Since these costs have been decreasing over the last decades due to more advanced hardware, Dai and Gao [11] state that it is nowadays only consequential to apply ML to the diagnostics process of CBM.

ML can be divided into two major types: supervised and unsupervised learning [13]. In addition, several hybrid types published in literature are beyond the scope of this paper due to their rare usage for FDD in the past [15]. The basic types can be distinguished based on labeled or unlabeled training data [11]. While unlabeled input data does not contain any information about the machine state, labeled data has been preprocessed to indicate whether the machine is in normal or faulty state during observation time. If multiple faults are classified, the type of fault is indicated [11].

Supervised Learning. Supervised learning requires labeled input data [11]. This paradigm tries to find a generalized description that maps all possible data points to the appropriate label. To achieve this, the technique tries to find patterns and rules that describe correlations between certain characteristics of the data and faults [11].

As mentioned earlier, diagnostics tasks are classification problems. Classification decides based on previously learned examples to which of N given classes a data point belongs. In discrete classification problems, each data point belongs to exactly one class [13].

Supervised ML can be split into two groups: classification and regression ([13, 16, 17]). As opposed to classification, regression aims at predicting one variable from known values of at least one other variable using a mathematical function [18]. Since in FDD the task is to decide if a component is in healthy or faulty state a decision boundary separating the states has to be identified. To generate boundaries from historic data, ML classification techniques can be used [13]. Regression, on the other hand, can be used for prognostics, however is not suited for diagnostics [19].

Unsupervised Learning. Although, supervised learning being one of the most frequently applied ML types [20], in many cases it is difficult to obtain labels [13]. Often, at least a certain amount of labels has to be created by hand, involving high expenses [21]. Consequently, another type of ML can be applied in an unsupervised context, where it learns on unlabeled input data. In this case, the classification tags, which represent the state of the data point, are missing. In unsupervised learning, these tags are then derived by the ML technique by identifying similarities within the data [11]. Similar data points are identified as belonging to the same category, which is also known as clustering [13].

2.3 Machine Learning Techniques Applied to Condition-Based Maintenance

To identify commonly used ML techniques for fault diagnostics, a structured literature review to has been conducted according to Vom Brocke et al. [22]. Results have shown that several ML techniques have been applied to fault diagnostics. Nevertheless, this has rarely been done in a structured manner analyzing the problem characteristics and comparing a larger number of algorithms at the same time. In this paper, a total number of eight techniques will be considered further, because they have been applied frequently or are interesting due to their potential future impact.

Artificial Neural Networks (ANN) and especially Support Vector Machines (SVM) have been identified as the most popular ML techniques during the last decade for CBM. Further identified techniques include Decision Trees (DT), the combined classifier Random Forest (RF) and K-Nearest Neighbor (K-NN). Even though K-NN has been employed only occasionally, it is considered here due to its simplicity and thus as baseline technique to enable comparison.

Three techniques from the area of unsupervised learning are also taken into consideration. The first, K-Means has not been used often but is still an integral part of unsupervised learning. A second more sophisticated method, the Self-Organizing Map (SOM), based on an ANN, will be taken into consideration as well. Finally, the Deep

Neural Network (DNN) developed in the last decade with increasing application in the last years can be applied for both supervised and unsupervised learning.

3 Requirement Catalogue from Condition-Based Maintenance for Assessment of Machine Learning Techniques

In order to be able to assess ML techniques, requirements, used to compare the techniques, need to be investigated. Requirements and their ratings for specific techniques are identified by conducting a second structured literature research.

A Requirements Catalogue for CBM. Venkatasubramanian [10] defines a catalog containing “Desirable features of a diagnostic system” (p. 1256). He summarizes general requirements of different kinds of diagnostics systems. As explained previously, ML is only one approach used for FDD. That is why the characteristics described in the following, are only to a certain extent valid for ML. The reason to apply ML in the context of CBM can be described as the need for a reliable, robust and efficient classifier that is able to identify multiple failures in complex machinery systems [12]. It aims to achieve the highest possible accuracy. This is the number of correct predictions, divided by the total number of predictions [20]. Accuracy is one of the measures used later to decide on the best ML technique for an application. In the following, out of the ten requirements presented by Venkatasubramanian [10], only the ones usable in the context of ML for FDD are described:

1. **Robustness:** During the application of a ML classifier, it is important to create a robust classifier, which is insensitive to various uncertainties and noise in the data. This can be achieved by tuning the threshold conservatively. This, on the other hand, leads to the trade-off between quick detection and diagnosis.
2. **Novelty Identifiability:** Since the ML classifier is designed to decide in which condition the machine is, the classifier needs to be able to distinguish between three cases: ‘normal operation’, ‘known fault mode’ and ‘unknown fault mode’. In the case of an ‘unknown fault mode’, the classifier is required to identify the novelty of this fault mode in comparison to the known fault modes.
3. **Classification Error Estimate:** In order to gain the users confidence in the classifier’s accuracy, the classifier is required to provide an a priori estimate of the error that might occur during classification. Thereby, the user can verify the current level of confidence a ML classifier has.
4. **Adaptability:** After the initial development of the model, the modeled process is subject to changes caused, for instance, by variations in external inputs. Consequently, the classifier should be adaptable to possible changes. This requirement, known as concept drift, is usually not met by basic techniques. However nowadays, many variants of techniques exist, which are able to adapt to new situations and learn when new data becomes available [12, 23]. Thus, in Table 1 a circle in the adaptability row describes the availability of extensions for the techniques.
5. **Explanation Facility:** The classifier should be able to demonstrate why it choose which fault class and how the fault has evolved.

6. Storage and Computational Requirements: When using ML classifiers for real-time classification, a tradeoff between computational complexity and high storage requirements exists. A reasonable balance between these two aspects is necessary.

The excluded requirements in the above list are Quick Detection and Diagnosis, Isolability, Modeling Requirements and Multiple Fault Identifiability. Regarding the first requirement, every ML technique's sensitivity is adjustable depending on the parameters that the user sets prior to training. Isolability in the context of ML reduces tolerance to noise of a technique, which is already measured by the Robustness (1) requirement.

Furthermore, Modeling Requirements are less relevant in the context of ML, because most of the intelligence behind the application of ML has shifted towards the feature extraction process. This process only depends on the data preparation and not on the specific ML technique that is used [24]. Multiple Fault Identifiability depends on the training set and a techniques ability to perform incremental learning. In this paper it is assumed, that neither a complete training set exists nor special incremental learning versions of the techniques are employed. A technique's ability to perform incremental learning is therefore assessed by the Adaptability (4) requirement.

Assessment of ML Techniques. To assess ML techniques, on the one hand the requirements described in the previous section will be used. On the other hand, general requirements for ML, which are valid in every application area of ML are taken from Kotsiantis [21]. They were obtained from previous studies. Because they are important for the assessment of ML techniques, the majority of these requirements are included in Table 1. The assessment of the ML techniques with regard to these requirements is performed considering 26 different publications, identified during the second structured literature review. The textual description and valuation of these publications have been used for the assessment of the ML techniques presented in Table 1 in form of a score from '--' to '++'.

In Table 1 the rating of most categories, except categories 2, 10 and 12 for the techniques ANN, SVM, DT and K-NN, are taken from Kotsiantis [21]. In case of the accuracy of the ANN, the rating was altered due to further advances in the development of ANNs during the last decade. Furthermore, the SVM's ability to deal with overfitting is updated according to newer results from the literature. The missing 60% of the ratings were obtained by additionally reviewing recent literature that contained a description and assessment of those techniques. Reviewed literature included well-accepted books on the area of ML and AI as well as specialized publications dealing with certain techniques. The score is then graded based on the author's descriptions. In the end, the general rating per technique is calculated by averaging the ratings of the twelve categories. All ratings are weighted with a factor of 1. The category accuracy, however, is weighted with the factor 5, to highlight its importance for the assessment of ML techniques. In case a different weighting of criteria is applicable with regard to the specific problem characteristics, the rating can be adapted accordingly. The final rating is obtained by adding all ratings together according to their weights and rounding the final decimal number to the next higher natural number. In the remainder of this chapter, the

most important pros and cons of the presented ML techniques are discussed based on Table 1.

Table 1. Scoring of the eight ML techniques based on the presented criteria: ‘--’ represents the worst (0), ‘o’ the neutral (3) and ‘++’ the best (5) rating; accuracy is weighted with a factor of 5, while all other categories are weighted with a factor of 1.

ML Technique:	ANN	SVM	DT	RF	K-NN	K-Means	SOM	DNN
<i>Properties of Results</i>								
1. Accuracy in general (5x)	+	+	-	++	-	-	o	++
2. Classification Error Estimate	++	++	o	++	+	+	++	+
3. Speed of classification	+	+	+	+	-	+	+	+
4. Transparency of classification – Explanation Facility	--	--	++	++	-	o	++	o
<i>Technique Abilities</i>								
5. Speed of learning with respect to number of attributes and instances	-	--	o	+	++	+	-	o
6. General data type rating:	+	o	++	++	+	-	o	o
6.a Dealing with discrete/	o	-	+	+	o	o	o	o
6.b binary/	+	+	+	+	+	-	o	+
6.c continuous attributes	+	+	+	+	+	+	+	o
7. General robustness rating:	-	o	-	+	--	-	+	+
7.a Tolerance to noise/	-	-	-	o	--	-	+	+
7.b missing values/	--	-	o	+	--	--	+	+
7.c irrelevant attributes/	--	+	o	+	-	-	+	+
7.d correlated attributes	o	o	-	++	-	+	o	-
8. Attempts for incremental learning – Adaptability	o	o	-	o	+	o	-	-
9. Dealing with overfitting	--	+	-	++	o	-	--	+
10. Novelty Identifiability	+	-	--	o	+	-	+	o
<i>Model Handling</i>								
11. Model Parameter Handling	-	-	+	o	+	+	-	--
12. Storage and Computation	o	-	+	+	--	+	o	-
General rating	+	+	o	++	o	o	+	+

Even though SVMs are in general slow during training, there are good classifiers, achieving over all high accuracy. Similar to ANNs, SVMs have a very poor Explanation Facility, since its choices of features are sometimes hard to trace [21]. In general, only

RFs and DNNs achieve better accuracy. SVMs, RFs and DNNs have the ability to deal with overfitting, while they also have acceptable Robustness. In general, RFs achieve the highest overall rating, because they lack major drawbacks. DNNs consist of several parameters that need to be set and optimized iteratively leading to a long learning process. For this reason, DNN performs poor in the Parameter Handling category. Conclusively, a technique superior in every aspect was not identified. In order to select appropriate techniques for a given data set, it is necessary to identify the properties of the data, such as the amount of noise and the type of attributes to select the best technique.

4 Validation

In order to demonstrate the benefits of the previously established rating of the eight ML techniques, appropriate techniques are selected for an application case and later applied to real world data. To be able to make an appropriate selection of techniques it is necessary to identify the problem characteristics of the application case with regard to the above-identified criteria. In this paper, this process is demonstrated based on an application case in the in-vitro diagnostics (IVD) industry.

Requirements Identification. Requirements for the application in the IVD industry needed to be obtained to proof the viability of the presented approach. This process is carried out in cooperation with STRATEC Biomedical AG, an IVD instruments manufacturer. After data acquisition, it has to be evaluated if ML can be used to produce failure predictions. When using ML, a core problem is the identification of an appropriate technique to give accurate predictions using condition-monitoring data. Therefore, potential requirements from STRATEC are identified first. For this reason, a self-completion questionnaire is designed and handed to leading employees of STRATEC from the area of (Project) Management, System Integration and Engineering. The questionnaire covers questions about STRATEC's maintenance activities. The participants are also asked to rate the importance of accuracy and explanation facility of a classifier. To determine whether false positives or true negatives are more important, the two cases are described and the participants are asked to select the less desirable case (based e.g. on costs or customer satisfaction).

In total, nine requirements could be identified from the results of the questionnaire, while the four most important of the nine requirements are presented in the following:

- I. Labels for the training sets are mostly available.
- II. False alarms need to be avoided if possible.
- III. High accuracy is important.
- IV. In the context of CBM in IVD the classifiers' Explanation Facility is less relevant.

After having established the pros and cons of the eight ML techniques, based on literature using the requirements from IVD, it is possible to select the most promising techniques for application in IVD. As explained in (I), labels for training are available, which enables supervised learning. For this reason, the two unsupervised clustering techniques (SOM, K-Means) have been excluded from the further analysis.

The questionnaire indicated that the accuracy of a classifier is important for the considered company (III), for this reason SVM, RF and DNN are selected. A second reason for the selection of these three techniques is that they are able to achieve above average general rating. The transparency of classification (IV) is not a very important factor in the context of CBM in IVD as long as the ML classifier provides acceptable performance. Therefore, techniques like ANN or SVM can be applied to the problem, which would have been impossible if high transparency was required. The two remaining methods from supervised learning are K-NN and DT. In this case study, DT is additionally selected to compare its performance to the combined classifier RF.

Evaluation of Techniques. To evaluate the techniques ANN, SVM, RF, DT and DNN, a test setup was created in close cooperation with STRATEC. Data was acquired from a reliability run of ten peristaltic pumps. Potential signals to monitor the condition of the pump were evaluated. The most promising approach was monitoring of motor currents. Based on the motor currents, the failures of tubing inside the peristaltic pump should be analyzed and predicted. Therefore, the motor signals were preprocessed using feature extraction. Extracted features, which were used for model development, include the range of the motor current over time, as well as the standard deviation and kurtosis. Additionally the skewness of the data as well as a crest factor were determined and added to the data set. Subsequently, in this binary classification problem, class 0 represents the ‘no fault’ class and class 1 represents the ‘faulty’ class. Since 10 pumps were tested under the same conditions, 9 out of 10 data sets were combined to form the training set (approx. 13000 rows). Pump 10 was excluded from the training set and was therefore used to test the classifier (approx. 1500 rows).

After data preparation, the selected ML techniques were applied on the training data using processes inside RapidMiner. The models were then tested on the test data to acquire the classification results. RapidMiner implemented four out of five techniques while the ‘Deep Learning’ Module within RapidMiner was designated to act as two layer DNN using stochastic gradient descent as optimizer.

Most parameters were kept with their default value. However, some selected parameters were altered. The LibSVM module with a svm type of nu-SVC and a sigmoid kernel was used as SVM. An ANN was configured using one hidden layer size with a size of 5, 500 training cycles and a learning rate of 0.35. The RF parameters were set to gain_ratio as the split criterion, a number of 21 trees and a maximum depth of 20. A decision tree was built, using information_gain as the split criterion and a pruning confidence of 0.1. Deep Learning was set to non-reproducible, using rectifier as the activation function and a number of 10 training epochs.

In order to assess the performance of the ML techniques, the performance measurements accuracy and recall are considered. While accuracy reflects the total number of misclassifications (independent of the class), the weighted recall and precision penalize false positive alarms by weighting false positives with a factor of 2. This is imposed by requirement (II) obtained earlier.

The results presented in Fig. 1 meet the expectations, given the assessment of the eight techniques in Sect. 3. The RF achieves outstanding performance producing only one misclassified example. The ANN performs very well in this problem. This is because

an ANN can yield very good results independent of the problem. Its scoring presented in Sect. 3 was mostly lowered by its low robustness, the danger of overfitting and its low transparency. None of these drawbacks is of importance for the problem at hand. As presented in Fig. 1 DT is inferior to all other techniques, although this was to be expected. In general, it can be noted that all techniques have been optimized to achieve a very high recall of class 0 (no fault). Thereby, fewer data points are classified as ‘faulty’, even though some of them are actually ‘faulty’. The ‘Deep Learning’ module yielded mixed results during application, whereas an acceptable solution could be obtained after several iterations. However, it performs worse than the RF, SVM and ANN in this problem.




Fig. 1. Comparison of classification results of the five classification techniques.

Subsequently, the best techniques for application in IVD are selected. Since all methods except the DT have a very high recall in class 0 (no fault), the four methods of SVM, ANN, RF and DNN are applicable to CBM in IVD. In order to combine the advantages of each technique, an ensemble of these classifiers is proposed in this context.

5 Conclusion

In this paper, eight well-known ML techniques have been assessed according to different criteria obtained from literature. To display the results, they have been combined in Table 1, which allows for easy reuse of the ratings for different tasks. Table 1 extends previous work done by Kotsiantis [20] by investigating further sources (in combination with an update of relevant aspects of the already assessed techniques) as well as adding additional requirements relevant in the context of fault diagnostics. Furthermore, four additional techniques have been included in the assessment. The eight techniques have been identified performing a literature research beforehand and adjusted to fit the specific requirements for fault diagnostics. These techniques include both supervised techniques (ANN, SVM, DT, RF, K-NN, K-Means) as well as unsupervised techniques (SOM, DNN). It was shown that, with regard to the identified criteria, Random Forests exhibit the best overall performance. However, Neural Networks, Support Vector Machines,

Self-Organizing Maps and Deep Neural Network score well. Depending on the specific application, different criteria could be weighted differently leading to other results. The application of the developed rating is demonstrated using a case study from the In-Vitro Diagnostics industry. During the case study, requirements for CBM from IVD have been obtained, using a self-completion questionnaire. These requirements were used to select five supervised techniques. The techniques are applied to a real-world data set from IVD. Their performance was presented and compared based on quantitative measures.

Our approach can assist other researchers at selecting the best-suited ML technique(s) in other application areas of CBM. This is beneficial, even though in this paper the ratings were limited to only eight ML techniques, while there are plenty of techniques available. Generating additional ratings for those techniques could be part of future research. Furthermore, there might be more criteria available and usable to distinguish ML techniques. In summary, while our rating is very useful in the context of CBM, there might be other fitting criteria for other ML application areas.

References

1. Mahantesh, N., Aditya, P., Kumar, U.: Integrated machine health monitoring: a knowledge based approach. *Int. J. Syst. Assur. Eng. Manage.* **5**(3), 371–382 (2014)
2. Shin, J.H., Jun, H.B.: On condition based maintenance policy. *J. Comput. Des. Eng.* **2**(2), 119–127 (2015)
3. Venkatasubramanian, V., Rengaswamy, R., Yin, K., Kavuri, S.N.: A review of process fault detection and diagnosis part I: quantitative model-based methods. *Comput. Chem. Eng.* **27**(3), 293–311 (2003)
4. Jardine, A.K.S., Lin, D., Banjevic, D.: A review on machinery diagnostics and prognostics implementing condition-based maintenance. *Mech. Syst. Signal Process.* **20**(7), 1483–1510 (2006)
5. Peng, Y., Dong, M., Zuo, M.: Current status of machine prognostics in condition-based maintenance: a review. *Int. J. Adv. Manuf. Technol.* **50**(1–4), 297–313 (2010)
6. Lee, J., Wu, F., Zhao, W., Ghaffari, M., Liao, L., Siegel, D.: Prognostics and health management design for rotary machinery systems – reviews, methodology and applications. *Mech. Syst. Signal Process.* **42**(1–2), 314–334 (2014)
7. Gao, Z., Cecati, C., Ding, S.: A survey of fault diagnosis and fault-tolerant techniques-part I: fault diagnosis with model-based and signal-based approaches. *IEEE Trans. Industr. Electron.* **62**(6), 3757–3767 (2015)
8. Vachtsevanos, G., Lewis, F., Roemer, M., Hess, A., Wu, B.: *Intelligent Fault Diagnosis and Prognosis for Engineering Systems*. Wiley, Hoboken (2006)
9. Schwabacher, M., Goebel, K.: A survey of artificial intelligence for prognostics. The intelligence report. In: *Association for the Advancement of Artificial Intelligence AAAI Fall Symposium 2007*, pp. 107–114. AAAI Press, Arlington (2007)
10. Venkatasubramanian, V.: Prognostic and diagnostic monitoring of complex systems for product lifecycle management: challenges and opportunities. *Comput. Chem. Eng.* **29**(6), 1253–1263 (2005)
11. Dai, X., Gao, Z.: From model, signal to knowledge: a data-driven perspective of fault detection and diagnosis. *IEEE Trans. Industr. Inf.* **9**(4), 2226–2238 (2013)
12. Wu, F., Wang, T., Lee, J.: An online adaptive condition based maintenance method for mechanical systems. *Mech. Syst. Signal Process.* **24**(8), 2985–2995 (2010)

13. Marsland, S.: *Machine Learning: An Algorithmic Perspective*, 2nd edn. Chapman and Hall/CRC Press, Boca Raton (2014)
14. ISO/IEC 2382.: *Information Technology – Vocabulary*. International Organization for Standardization, Geneva (2015)
15. Zhao, X., Li, M., Xu, J., Song, G.: Multi-class semi-supervised learning in machine condition monitoring. In: *2009 International Conference on Information Engineering and Computer Science*, pp. 1–4. IEEE Press, Wuhan (2009)
16. Bishop, C.M.: *Pattern Recognition and Machine Learning*. Springer, New York (2006)
17. Chapelle, O., Scholkopf, B., Zien, A.: Introduction to semi-supervised learning. In: Chapelle, O., Scholkopf, B., Zien, A. (eds.) *Semi-Supervised Learning*, pp. 1–12. The MIT Press, Cambridge (2006)
18. Ahmad, R., Kamaruddin, S.: A review of condition-based maintenance decision-making. *Eur. J. Industr. Eng.* **6**(5), 519–541 (2012)
19. Susto, G.A., Schirru, A., Pampuri, S., McLoone, S., Beghi, A.: Machine learning for predictive maintenance: a multiple classifier approach. *IEEE Trans. Industr. Inf.* **11**(3), 812–820 (2015)
20. Kotsiantis, S.B.: Supervised machine learning: a review of classification techniques. *Informatica* **31**(3), 249–268 (2007)
21. Chandola, V., Banerjee, V., Kumar, V.: Anomaly detection: a survey. *ACM Comput. Surv.* **41**(3), 1–58 (2009)
22. Vom Brocke, J., Simons, A., Niehaves, B., Riemer, K., Plattfaut, R., Clevén, A.: Reconstructing the giant: on the importance of rigour in documenting the literature search process. In: *ECIS*, vol. 9, pp. 2206–2217 (2009)
23. Widmer, G., Kubat, M.: Learning in the presence of concept drift and hidden contexts. *Mach. Learn.* **23**(1), 69–101 (1996)
24. Arel, I., Rose, D.C., Karnowski, T.P.: Deep machine learning – a new frontier in artificial intelligence research. *IEEE Comput. Intell. Mag.* **5**(4), 13–18 (2010)



A Comparative Evaluation of Log-Based Process Performance Analysis Techniques

Fredrik Milani and Fabrizio M. Maggi^(✉)

University of Tartu, Tartu, Estonia
milani@ut.ee, f.m.maggi@ut.ee

Abstract. Process mining has gained traction over the past decade and an impressive body of research has resulted in the introduction of a variety of process mining approaches measuring process performance. Having this set of techniques available, organizations might find it difficult to identify which approach is best suited considering context, performance indicator, and data availability. In light of this challenge, this paper aims at introducing a framework for categorizing and selecting performance analysis approaches based on existing research. We start from a systematic literature review for identifying the existing works discussing how to measure process performance based on information retrieved from event logs. Then, the proposed framework is built starting from the information retrieved from these studies taking into consideration different aspects of performance analysis.

Keywords: Process mining · Performance analysis
Evaluation framework

1 Introduction

Businesses are at a turning point where they have to incorporate digitalization or fade away. Digital technologies continue to set their transformative marks on virtually all industry domains and have allowed the expansion of businesses to markets previously inaccessible. The forces of innovation and creativity have enabled young businesses to challenge incumbents in practically every sector. However, one thing has not changed. Businesses will always seek to improve their processes because “every good process eventually becomes a bad process” [16]. This is even more relevant in a fast-changing digital era.

The first step to process improvement is to understand where processes can be improved. In the past, given the lack of data availability and high cost of data processing, performance analysis methods identified improvement opportunities based on manual analysis, and at times combined with random sampling (e.g., six sigma [28]). Relying on such manually driven methods, process analysts assessed the performance of processes so to find opportunities for improvement. Today, much of the data is captured digitally and, over the past decade, analysis of large sets of data has improved remarkably. No longer are businesses restricted

to select the most prioritized processes, limit the scope, or confine the selection of data due to limitations of time-consuming analysis or tools. In addition, the accessibility to open source tools has never been easier, in particular for data driven analysis of business processes. Therefore, in a digital era, businesses cannot hope to survive with manually driven methods. The process analysis must also be digitally transformed by tapping into data-driven analysis methods.

One group of techniques for data driven performance analysis uses event logs of processes to assess performance. Indeed, nowadays, business processes are often supported by IT systems that log their execution. For instance, an order-to-invoice process might include activities such as REGISTER, VALIDATE, APPROVE, FILL ORDER and SEND INVOICE. Each order has a unique id and every activity is recorded in the event log with information about the time when it was executed (timestamp) and other additional data such as the resource (person or device) that executed the activity. As such, the process is inherently captured in the log. With process mining techniques [2], the performance of such processes can be assessed and analyzed in great detail based on event logs.

The body of research and tools within the field of process mining has grown significantly during this decade. However, the availability of tools and approaches developed for specific aspects of process performance does not make it easier for businesses to employ them. In fact, it poses a challenge. There is no way for businesses to easily get an overview of what performance indicators can be measured, what input data is required for such analysis, or what industry specific implementations are available. In light of this context, we propose a framework for the selection of log-based performance analysis techniques. We do so by conducting a systematic literature review to identify the body of existing work. We analyze the results and focus on identifying existing process performance indicators, required input data, and approaches available. Based on the results, we build a framework for the selection of suitable performance analysis approaches.

The structure of this paper is as follows. Section 2 summarizes the research protocol for the systematic survey. In Sect. 3, the research questions are discussed and the framework is presented in Sect. 4. Finally, Sect. 5 concludes the paper.

2 Systematic Literature Review

In this section, we summarize how the systematic literature review was conducted. The review protocol specifies research questions, search protocol, inclusion and exclusion criteria, and data extraction. The review protocol predominantly follows the guidelines provided by Kitchenham [20]. The objectives of this paper are to review the current academic research on performance analysis techniques based on logs and build a framework for categorizing them. To this end, the overarching research question of “what is the body of relevant academic research within the field of process performance analysis?” has been decomposed into three sub-questions:

- **RQ1:** *What aspects of process performance do existing techniques consider?*
This research question aims at identifying the aspects that can be measured in regards to process performance.

- **RQ2:** *What input data is required for measuring process performance?* For performance analysis, it is important that the “right” set of data is captured. To this end, it is important to understand what kind of data is required as input for process performance analysis techniques.
- **RQ3:** *What are the main approaches/algorithms and tools available for analysis of process performance?* The final research question aims at capturing the various methods that can be applied for performance analysis.

To find relevant studies, we sought studies within the domain of “process mining”. However, as process mining covers many aspects of business process analysis, such as process discovery [1], we included “performance” to focus the search. The boolean search string (“process mining” AND “performance”) was used. The search was applied to Web of Science (WoS) and Scopus databases. These electronic databases were selected as they are the primary databases for computer science research. The search was conducted in January 2018 and resulted in a total of 330 studies from Scopus and 194 from WoS. After having removed duplicates, 349 studies remained. The first filtering was aimed at removing studies that were clearly out of scope (based on title), shorter than 6 pages, not accessible, or not written in English. The abstract and introduction of the remaining 101 studies were examined. Peer reviewed studies introducing or extending an existing approach for performance analysis, or directly dealing with performance of a process were included. After this filtering, the list of studies was reduced to 48. These studies were examined in full following the inclusion criteria of being within the field of log-based performance analysis (IC1), proposing a new approach (IC2) or extending an existing approach (IC3) for measuring process performance, and applying the presented method(s) to an industrial case study (IC4).

The final list of studies, following the above criteria, constitutes of 32 studies. For each study in the final list, standard metadata was extracted. In order to address the first research question, data about performance category, metric, and unit of measurement was extracted. For the second research question, information about what input data is required to do the analysis was extracted. Finally, for the third research question, information about tools used and underlying method was extracted.

3 Results

In this section, we examine the final list of studies from the perspective of the defined research questions. The first being the different aspects that can be assessed, followed by the data input required and available tools and methods.

3.1 Aspects of Process Performance (RQ1)

The first research question concerns what aspects of process performance existing techniques measure. Not surprisingly, we found that the majority of the analyzed studies include analysis of the time perspective. This is perhaps the most basic performance aspect included in all mature performance measuring tools.

Time. The time aspect can be divided into four categories. These are process, fragment, activity and waiting duration (see Fig. 1).

Process Duration. Process duration is the time distance between the start event of a process and the end event. Several techniques measure process duration. For instance, [34] examines the process duration of a peer-review process to identify bottlenecks. Similarly, Engel et al. [11] use electronic exchange messages to analyze the duration of inter-organizational processes. Different aspects of process duration can also be analyzed. For instance, [23] examines the influence of contextual factors (such as weekday or season) on process duration. Ballambettu et al. [5] propose a method for identifying key differences of process variants that could affect process duration. Suriadi et al. [39] look at the processes of the emergency departments at four different hospitals. They compare these processes and their process duration to identify differences. Piessens et al., [32] recognize that some event logs contain advanced constructs such as cancelations, multiple concurrent instances, or advanced synchronization constructs. They use these constructs to gain accurate assessment of process duration.




Fig. 1. Aspects of time performance

Fragment Duration. Fragment duration considers the time required to complete a fragment (a set of activities) of a process. Wang et al. [41] propose a framework for applying process mining in logistics and analyze process fragments of a Chinese bulk port process. They identify the most time-consuming fragments of the process and, using the fragment durations, they categorize cases containing those fragment to give insights on their performance.

Activity Duration. Activity duration considers the duration of an activity. Activity duration analysis is also very common in performance analysis. For instance, [7] applies existing process mining techniques to analyze the activity duration of a Korean hospital event log. A similar analysis is conducted by [24] on a Dutch hospital log. Leyer [23] measures the impact of contextual factors on activity duration. Activity time is analyzed in a two-step method combining process mining and statistical methods. In a similar vein, Hompes et al. [17] use statistical methods to analyze the effect of context on a set of key process performance indicators at the activity level. Activity duration analysis commonly takes an aggregated view-point, considering, for instance, the average time of all executed instances of each

activity. However, a process have variants where activity durations vary across different variants. To address variability in activity duration based on variants, in [5], the authors propose a method that allows for identifying key differences of activity duration across process variants.

Waiting Duration. Waiting time in processes is one of the main wastes [10] and, as such, it has to be reduced for process improvement. As it is one of the main approaches to improve processes, waiting time analysis is the focus of many techniques. Jaisook and Premchaiswadi [19] investigate hospital logs to examine the average duration a patient spends waiting in a private hospital. The authors do so by using the built-in functionality of Disco [14]. Similarly, Perimal-Lewis et al. [31] rely on Disco to examine the processes of an emergency department. In so doing, they apply process mining to identify deviating activities in regards to waiting duration. The results highlight bottlenecks in the performance of the processes. Park et al. [29] propose a framework for analyzing block manufacturing processes by assessing the total waiting time. In a similar manner, in [34], the authors examine the total waiting time of a peer-review process. In [5, 6], the authors present a framework for analyzing similar processes across several installations. They propose a method to analyze a collection of logs from different performance perspectives, one of which is waiting duration. An aspect of waiting duration is delay analysis. Delays refer to cases where the completion time is later than the planned completion time. Senderovich et al. [38] analyze a process log from the perspective of operational deviations resulting in tardiness (delays) from a process duration perspective. Park et al. [30] analyze delays in a make-to-order manufacturing firm. They define two delay indicators, activity and processing delay and found that some delays can be explained by seasonality.

Resources. The performance of human resources is another aspect of process performance that is often analyzed. Pika et al. [33] introduce an extensive framework for analysis of human resources from different perspectives. Their framework measures with the aid of time series analysis, resource utilization and productivity. Workload has been recognized as affecting resource performance as discussed in [25]. Here, the authors explore the effect of workload on service times based on historic data and by using regression analysis. A similar metric is used by [30] when analyzing manufacturing processes. In [18], Huang et al. present an approach for measuring resource behavior from four perspectives, i.e., preference, availability, competence and cooperation. Resources can also be non-human such as materials. In analyzing a block manufacturing process, Park et al. [29] consider materials (welding length) for performance analysis.

Quality. We also identified quality as a performance perspective. Quality can be divided into internal and external. Internal quality regards the conformance of the process outcome to internally defined targets, whereas external quality refers to customers' satisfaction with the process outcome [10]. Internal quality analysis has been conducted by Arpasat et al. [4] who analyze the reasons why too many

attempts were required to solve a problem in a bank customer-service process. They apply Disco to identify the causes for inappropriate (not successful) interventions. External quality can be based on the analysis of the complaints received. Wongvigran and Premchaiswadi [43] analyze a call-center log by considering the number of complaints to identify teams receiving most complaints.

3.2 What Input Data (RQ2)

Automated performance analysis requires logs capturing executed events. In order to apply process mining techniques, the minimum requirement on data captured by event logs are “case id”, “activity”, and “timestamp” [2]. However, for performance analysis, additional data is required depending on what kind of analysis is to be conducted. The time performance of a process, be it the process, waiting, activity, or fragment duration, is commonly measured as maximum, minimum, mean, or average duration [32, 34]. A mere timestamp is sufficient when considering process or fragment duration as the duration is calculated by using the timestamps of the first and the last activity (of the process or fragment). However, if activity and waiting duration are to be measured, it is necessary to have the start and end time (timestamp) of each activity in the log [6, 7, 17, 24]. Indeed, logs might only include one timestamp such as for activity completion. Activity and waiting duration analysis is not possible on such logs. However, in these cases, it is possible to estimate average activity and waiting times using probabilistic methods [27].

Resources are measured by considering the “performer” representing the human resource and/or “materials”. Materials refer to the amount of a particular type of materials used for performing an activity. For such purposes, the input log must hold data on performers and/or amounts (quantity or costs). Semi-structured business processes (where the execution of the process is not fully supported by a system) do not capture all interactions among actors (e.g., interactions with customers). Logs from such processes capture the information partially. Wombacher and Lacob [42] propose an approach to make such logs suitable for performance analysis.

Quality is either “internally” or “externally” induced. Internally induced measures commonly include binary categorization of the process outcome (desired or undesired) such as defects, errors, or delays [4]. Externally induced performance refers to the determination of quality from sources external to the process such as customer complaints [43]. The log must contain data that clearly marks each case with information about internally or externally induced measures. If the data exists outside the log, a pre-processing is required to enrich the log with the required quality attributes.

Table 1 depicts the data requirements (Y) and optional requirements (O) for process performance analysis. As can be seen, the minimum requirements are case id, activity, and timestamp. It follows naturally that the more data the log holds, the more advanced performance analysis can be made.

Table 1. Input data required for performance analysis

Attribute	Time				Resources		Quality	
	Process	Fragment	Activity	Waiting	Performer	Materials	Internal	External
Case Id	Y	Y	Y	Y	Y	Y	Y	Y
Activity	Y	Y	Y	Y	Y	Y	Y	Y
Timestamp	Y	Y	Y	Y	Y	Y	Y	Y
Activity Start Time	-	-	Y	Y	O	O	-	-
Activity End Time	-	-	Y	Y	O	O	-	-
Quality Tag	-	-	-	-	-	-	Y	Y
Performer	-	-	-	-	Y	O	-	-
Materials	-	-	-	-	O	Y	-	-

3.3 Approaches and Tools (RQ3)

There is a range of techniques to extract and analyze process performance characteristics (incl. performance measures) from event logs. For example, de Leoni et al. [21,22] propose a framework to extract process performance characteristics from event logs and to correlate them in order to discriminate, for example, between the performance of cases that lead to “positive” outcomes versus “negative” outcomes. In [33], the authors present an extensible framework for extracting knowledge from event logs about the behavior of a human resource and for analyzing the dynamics of this behavior over time.

Another group of works is aimed at understanding the influence of contextual factors on process performance. For example, in [17], the authors introduce a generic context-aware analysis framework that analyzes activity durations using multiple perspectives. In [35], Reijers et al. investigate whether the place where an actor works affects the performance of a business process. In [23], the authors present a methodological approach to identify the effect of contextual factors on business process performance in terms of processing time combining process mining techniques with statistical methods. This approach facilitates detecting impacted activities thus determining which activities within a business process are indeed dependent on the context. Close to the above studies are the ones presented in [5,6]. Here, starting from the observation that an organization might perform well for some clients and perform below par on others, the authors present a framework for analyzing operational event data of related processes across different clients to gain insights on process performance.

Another group of approaches is related to performance in collaborative processes. For example, in [11], the authors present the EDImine Framework for enabling the application of process mining techniques in the field of EDI-supported inter-organizational business processes, and for supporting inter-organizational performance evaluation using business information from EDI messages, event logs, and process models. In [15], Hachicha et al. present an analysis and assessment approach for collaborative business processes in SOA in order to maintain their performance in competitive markets.

Process performance has also been approached from the perspective of queuing theory. Senderovich et al. [36,37] propose a method to discover characteristics

of “work queues” from event logs at the level of an entire process or of individual activities. In [38], the authors target the analysis of resource-driven processes based on event logs. In particular, they focus on processes for which there exists a predefined assignment of activity instances to resources that execute activities. The goal is to decrease the tardiness and lower the flow time.

More advanced performance analysis techniques have been recently presented in [26,40]. In [26], the authors present a technique to understand how bottlenecks form and dissolve over time via the notion of Staged Process Flow. In [40], Suriadi et al. present a framework based on the concept of event interval. The framework allows for a systematic approach to sophisticated performance-related analysis (e.g., resource productivity trends, recurring working patterns of resources, waiting time distributions over time, and resource performance comparison), even with information-poor event logs.

Other studies overlay the performance measures on top of a process model by replaying the log on the process model [3,32] and calculating aggregate performance measures for each element in the process model during the replay. Techniques for enhancing the quality of performance analysis based on log replay have been proposed [9]. A related technique supported by contemporary performance analysis tools is log animation. Log animation displays in a movie-like fashion how cases circulate through the process model over time [8,9,13].

The analyzed studies mainly use ProM (15 studies) and Disco (6 studies). In the remaining studies, the authors developed their own applications.

4 Framework

In this section, we synthesize the above results in a framework aimed at assisting businesses to find the most suitable approach for performance analysis. Businesses, often not acquainted to the academic domain within this field, might find it challenging to navigate through the studies. As such, our framework might help in identifying the first steps.

The framework considers three types of techniques. Most studies aim at descriptive performance analysis of a single log. Concerning this type of techniques, we consider two aspects. The first aspect is derived from the first research question about performance perspectives. As such, the performance perspectives are time, resources, and quality. The second aspect refers to the data available in the input logs. Depending on what data is available, different types of the performance can be analyzed. Note that the log must include at least case id, activity, and timestamp (minimum required data). Some studies compare logs of similar processes from several sites or use logs pertaining to collaborative processes. Such approaches are more complex but might be highly relevant for some businesses. Finally, we noted case studies contextualizing performance analysis to a certain domain. As such case studies are also valuable for businesses, we include them in the framework. When combining all these techniques, we gain a framework as shown in Table 2.

A business seeking to conduct data-driven performance analysis, should first select the type of technique. Descriptive analysis will show the current state and

Table 2. Framework

Type	Input	Performance perspective	Resources	Quality
Descriptive performance analysis	Minimum Required Data	[3, 5, 6, 23, 32, 34] Process Duration [41] Fragment Duration	-	-
	Activity Start and End Time	[5, 7, 17, 23, 24] Activity Duration [7, 19, 24, 29, 31, 34] Waiting Duration [30, 36–38] Delay Duration	-	-
	Internal Quality	.	-	[4]
	External Quality	-	-	[43]
	Human Resources	-	[18, 25, 30, 33, 35]	-
Materials	-	[29]	-	
Type	Description			
Complex performance analysis	[21, 22] Framework to extract process characteristics from event logs discriminating between positive and negative cases			
	[5, 6] Comparing waiting duration of similar process in different installations			
	[11, 15] Collaborative Processes			
	[26] Evolution of performance over time			
[40] Framework for performance-related analysis with information-poor event logs				
Type	Domain			
Case study	[17, 23, 33] Banks			
	[7, 12, 19, 24, 39, 44] Healthcare Processes			
	[11, 29, 30] Manufacturing Processes			
	[41] Logistics			
[4, 25, 34, 43] Service Processes				

highlight cases and/or areas in the process where there are opportunities for improvement. For descriptive analysis, the minimum requirement is an input log capturing mandatory data (case id, activity, and timestamp). With this data, it is possible to perform process and fragment duration analysis. If the log contains timestamps for start and end of activities, it will be possible to conduct activity and waiting duration analysis. For delay analysis, it might be required to have scheduling data. For human resource performance analysis, the log must contain data about who performed which activity. However, resource does not need only to be human. For non-human resource analysis, the log must clearly show how much of the materials was used for each case or activity. For quality analysis, the log must also contain, for each case, data about if the case had a desired or undesired outcome. This might be in case of a defect or complaint made.

Our review reveals approaches used for complex performance analysis. Complex analysis covers comparative performance analysis between several installations, such as treatment processes at different hospitals or ERP systems installed at several client organizations. These approaches not only analyze the performance of each variant, but also compare them so to identify reasons for one being more efficient than the other. Another type of complex analysis is the one related to the performance of collaborative processes that can be inter- or intra-organizational. In addition, there are techniques to extract process characteristics from event logs to the aim of discriminating between positive and negative performance, techniques for the analysis of the evolution of process performance over time, and for sophisticated performance-related analysis with information-poor event logs.

Finally, the framework contains case studies from different industry domains. Most methods have validated their results on real-life industry logs and in so doing, also gained some insight that is specific for that industry. For instance, financial, healthcare, and manufacturing processes have been used to validate results. This will be valuable to businesses operating within the same industry or that have similar processes as those used for testing the results in the analyzed studies.

5 Conclusion

Business process performance analysis has been conducted for many decades to identify opportunities for process enhancement. In this light, performance analysis based on process mining techniques offer great value for businesses and our systematic literature review identifies tools for them to use. However, it might be difficult for businesses to navigate within this field. Therefore, we propose a framework to aid them in finding suitable methods. The framework considers the complexity of the analysis, performance perspectives, required input data, and tool availability. In particular, we show that performance is analyzed from time (process, fragment, activity, and waiting duration), resources, and quality aspects. Although process flexibility is also a performance indicator [10], currently there are no approaches for its analysis. Therefore, an important avenue

for future work in the process performance analysis field is the development of techniques for analyzing this performance perspective.

Acknowledgments. This project and research is supported by Archimedes Foundation and GoSwift OÜ under the Framework of Support for Applied Research in Smart Specialization Growth Areas.

References

1. van der Aalst, W.M.P.: Process discovery: an introduction. In: *Process Mining*, pp. 125–156. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-19345-3_5
2. van der Aalst, W.M.P.: *Process Mining - Data Science in Action*, 2nd edn. Springer (2016)
3. van der Aalst, W.M.P., Adriansyah, A., van Dongen, B.F.: Replaying history on process models for conformance checking and performance analysis. *Wiley Interdisc. Rew Data Min. Knowl. Discov.* **2**(2), 182–192 (2012)
4. Arpasat, P., Porouhan, P., Premchaiswadi, W.: Improvement of call center customer service in a Thai bank using Disco fuzzy mining algorithm. In: *ICT and Knowledge Engineering*, pp. 90–96 (2015)
5. Ballabettu, N.P., Suresh, M.A., Bose, R.P.J.C.: Analyzing process variants to understand differences in key performance indices. In: Dubois, E., Pohl, K. (eds.) *CAiSE 2017. LNCS*, vol. 10253, pp. 298–313. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-59536-8_19
6. Jagadeesh Chandra Bose, R.P., Gupta, A., Chander, D., Ramanath, A., Dasgupta, K.: Opportunities for process improvement: a cross-clientele analysis of event data using process mining. In: Barros, A., Grigori, D., Narendra, N.C., Dam, H.K. (eds.) *ICSOC 2015. LNCS*, vol. 9435, pp. 444–460. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-48616-0_31
7. Cho, M., Song, M., Yoo, S.: A systematic methodology for outpatient process analysis based on process mining. In: Ouyang, C., Jung, J.-Y. (eds.) *AP-BPM 2014. LNBIP*, vol. 181, pp. 31–42. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-08222-6_3
8. Conforti, R., Dumas, M., La Rosa, M., Maaradji, A., Nguyen, H., Ostovar, A., Raboczi, S.: Analysis of business process variants in Apromore. In: *BPM Demos*, pp. 16–20 (2015)
9. van Dongen, B.F., Adriansyah, A.: Process mining: fuzzy clustering and performance visualization. In: Rinderle-Ma, S., Sadiq, S., Leymann, F. (eds.) *BPM 2009. LNBIP*, vol. 43, pp. 158–169. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-12186-9_15
10. Dumas, M., La Rosa, M., Mendling, J., Reijers, H.A.: *Fundamentals of Business Process Management*, 2nd edn. Springer (2018)
11. Engel, R., Krathu, W., Zapletal, M., Pichler, C., Bose, R.J.C., van der Aalst, W.M.P., Werthner, H., Huemer, C.: Analyzing inter-organizational business processes. *Inf. Syst. e-Bus. Manag.* **14**(3), 577–612 (2016)
12. Ganesha, K., Supriya, K.V., Soundarya, M.: Analyzing the waiting time of patients in hospital by applying heuristics process miner. In: *ICICCT*, pp. 500–505 (2017)

13. Günther, C.W., van der Aalst, W.M.P.: Fuzzy mining – adaptive process simplification based on multi-perspective metrics. In: Alonso, G., Dadam, P., Rosemann, M. (eds.) BPM 2007. LNCS, vol. 4714, pp. 328–343. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-75183-0_24
14. Günther, C.W., Rozinat, A.: Disco: discover your processes. In: BPM Demos, pp. 40–44 (2012)
15. Hachicha, M., Fahad, M., Moalla, N., Ouzrout, Y.: Performance assessment architecture for collaborative business processes in BPM-SOA-based environment. *Data Knowl. Eng.* **105**, 73–89 (2016)
16. Hammer, M.: What is business process management? In: Brocke, J., Rosemann, M. (eds.) *Handbook on Business Process Management 1. International Handbooks on Information Systems*, pp. 3–16. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-00416-2_1
17. Hompes, B.F.A., Buijs, J.C.A.M., van der Aalst, W.M.P.: A generic framework for context-aware process performance analysis. In: Debruyne, C., et al. (eds.) *On the Move to Meaningful Internet Systems: OTM 2016 Conferences. OTM 2016. LNCS*, vol. 10033, pp. 300–317. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-48472-3_17
18. Huang, Z., Lu, X., Duan, H.: Resource behavior measure and application in business process management. *Expert Syst. Appl.* **39**(7), 6458–6468 (2012)
19. Jaisook, P., Premchaiswadi, W.: Time performance analysis of medical treatment processes by using Disco. In: *ICT and Knowledge Engineering*, pp. 110–115 (2015)
20. Kitchenham, B.: *Procedures for performing systematic reviews*. Keele University, Keele, UK, 33, pp. 1–26 (2004)
21. de Leoni, M., van der Aalst, W.M.P., Dees, M.: A general framework for correlating business process characteristics. In: Sadiq, S., Soffer, P., Völzer, H. (eds.) *BPM 2014. LNCS*, vol. 8659, pp. 250–266. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-10172-9_16
22. de Leoni, M., van der Aalst, W.M.P., Dees, M.: A general process mining framework for correlating, predicting and clustering dynamic behavior based on event logs. *Inf. Syst.* **56**, 235–257 (2016)
23. Leyer, M.: Towards A context-aware analysis of business process performance. In: *PACIS 2011: Quality Research in Pacific Asia*, p. 108 (2011)
24. Mans, R.S., Schonenberg, M., Song, M., van der Aalst, W.M.P., Bakker, P.J.: Application of process mining in healthcare—a case study in a Dutch hospital. In: *BIOSTEC*, pp. 425–438 (2008)
25. Nakatumba, J., van der Aalst, W.M.P.: Analyzing resource behavior using process mining. In: Rinderle-Ma, S., Sadiq, S., Leymann, F. (eds.) *BPM 2009. LNBIP*, vol. 43, pp. 69–80. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-12186-9_8
26. Nguyen, H., Dumas, M., ter Hofstede, A.H.M., La Rosa, M., Maggi, F.M.: Business process performance mining with staged process flows. In: Nurcan, S., Soffer, P., Bajec, M., Eder, J. (eds.) *CAiSE 2016. LNCS*, vol. 9694, pp. 167–185. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-39696-5_11
27. Nogayama, T., Takahashi, H.: Estimation of average latent waiting and service times of activities from event logs. In: Motahari-Nezhad, H.R., Recker, J., Weidlich, M. (eds.) *BPM 2015. LNCS*, vol. 9253, pp. 172–179. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-23063-4_11
28. Pande, P.S., Neuman, R.P., Cavanagh, R.R.: *The Six Sigma Way*. McGraw-Hill (2000)

29. Park, J., Lee, D., Zhu, J.: An integrated approach for ship block manufacturing process performance evaluation: case from a Korean shipbuilding company. *Int. J. Prod. Econ.* **156**, 214–222 (2014)
30. Park, M., Song, M., Baek, T.H., Son, S.Y., Ha, S.J., Cho, S.W.: Workload and delay analysis in manufacturing process using process mining. In: Bae, J., Suriadi, S., Wen, L. (eds.) *AP-BPM 2015*. LNBP, vol. 219, pp. 138–151. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-19509-4_11
31. Perimal-Lewis, L., Teubner, D., Hakendorf, P., Horwood, C.: Application of process mining to assess the data quality of routinely collected time-based performance data sourced from electronic health records by validating process conformance. *Health Inform. J.* **22**(4), 1017–1029 (2016)
32. Piessens, D., Wynn, M.T., Adams, M.J., van Dongen, B.F.: Performance analysis of business process models with advanced constructs. In: *Australasian Conference on Information Systems* (2010)
33. Pika, A., Wynn, M.T., Fidge, C.J., ter Hofstede, A.H.M., Leyer, M., van der Aalst, W.M.P.: An extensible framework for analysing resource behaviour using event logs. In: Jarke, M., et al. (eds.) *CAiSE 2014*. LNCS, vol. 8484, pp. 564–579. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-07881-6_38
34. Premchaiswadi, W., Porouhan, P.: Process modeling and bottleneck mining in online peer-review systems. *SpringerPlus* **4**(1), 441 (2015)
35. Reijers, H.A., Song, M., Jeong, B.: On the performance of workflow processes with distributed actors: does place matter? In: Alonso, G., Dadam, P., Rosemann, M. (eds.) *BPM 2007*. LNCS, vol. 4714, pp. 32–47. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-75183-0_3
36. Senderovich, A., Weidlich, M., Gal, A., Mandelbaum, A.: Queue mining – predicting delays in service processes. In: Jarke, M., et al. (eds.) *CAiSE 2014*. LNCS, vol. 8484, pp. 42–57. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-07881-6_4
37. Senderovich, A., Weidlich, M., Gal, A., Mandelbaum, A.: Queue mining for delay prediction in multi-class service processes. *Inf. Syst.* **53**, 278–295 (2015)
38. Senderovich, A., Weidlich, M., Yedidsion, L., Gal, A., Mandelbaum, A., Kadish, S., Bunnell, C.A.: Conformance checking and performance improvement in scheduled processes: a queueing-network perspective. *Inf. Sys.* **62**, 185–206 (2016)
39. Suriadi, S., Mans, R.S., Wynn, M.T., Partington, A., Karnon, J.: Measuring patient flow variations: a cross-organisational process mining approach. In: Ouyang, C., Jung, J.-Y. (eds.) *AP-BPM 2014*. LNBP, vol. 181, pp. 43–58. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-08222-6_4
40. Suriadi, S., Ouyang, C., van der Aalst, W.M.P., ter Hofstede, A.H.M.: Event interval analysis: why do processes take time? *Dec. Supp. Syst.* **79**, 77–98 (2015)
41. Wang, Y., Caron, F., Vanthienen, J., Huang, L., Guo, Y.: Acquiring logistics process intelligence: methodology and an application for a chinese bulk port. *Expert Syst. Appl.* **41**(1), 195–209 (2014)
42. Wombacher, A., Iacob, M.: Start time and duration distribution estimation in semi-structured processes. In: *SAC*, pp. 1403–1409 (2013)
43. Wongvigran, S., Premchaiswadi, W.: Analysis of call-center operational data using role hierarchy miner. In: *ICT and Knowledge Engineering*, pp. 142–146 (2015)
44. Yampaka, T., Chongstitvatana, P.: An application of process mining for queueing system in health service. In: *JCSSE*, pp. 1–6 (2016)



Blockchain for Business Applications: A Systematic Literature Review

Ioannis Konstantinidis¹(✉), Georgios Siaminos¹, Christos Timplalexis¹,
Panagiotis Zervas¹, Vassilios Peristeras¹, and Stefan Decker²

¹ International Hellenic University, Thessaloniki, Greece
{i.konstantinidis,g.siaminos,c.timplalexis,
p.zervas,v.peristeras}@ihu.edu.gr

² Fraunhofer Institute for Applied Information Technology FIT,
53754 Sankt Augustin, Germany
stefan.decker@fit.fraunhofer.de

Abstract. Blockchain technology is widely known as the technological basis on which bitcoin is built. This technology has created high expectations, as transactions of every kind are executed in a decentralized way, without the need of a trusted third-party. Blockchain real business applications are currently limited mostly to financial services but many R&D projects in companies and corporations try to amplify the areas of blockchain implementation. In this paper, we conduct a systematic survey with the aim of pointing out the areas in which blockchain technology applications and services are being developed both in the public and private sector. In the results, we discuss the disruptive effect that this technology could bring to various business sectors as well as the concerns regarding the development of the blockchain technology.

Keywords: Blockchain · Literature review · Business · Applications

1 Introduction

Blockchain is a digitized, decentralized public ledger intended to keep a record of every data transaction happening in its network. Every different user constitutes a network node and maintains a copy of the ledger. Each transaction on the blockchain database is verified by the users participating in the system, so a trusted third-party verification is not required.

In 2008, an unknown author, Nakamoto, wrote a paper about accomplishing non-reversible and cash-like transactions without the involvement of any third-party. This was blockchain's first use, the technology behind bitcoin cryptocurrency. The concept was quite simple [1, 2]. Suppose that user A wants to transfer money or data to user B. When this transaction happens, it is represented as a block which is transmitted to every node/user of the network. Then, the users have to verify if this transaction is valid. The users have to solve a puzzle in order to be the first to validate the transaction [3]. This puzzle demands the use of certain computational power. The puzzle solving procedure is called "mining" and the first miner who will find the solution gets a bitcoin reward, so miners are competing to be the fastest to solve the puzzle. The transaction is completed

when 51% of the users approve the provided solution. Then, the block of the transaction is added to the blockchain. The blockchain is a list of blocks that includes every single transaction that has ever been made [4, 5]. The blocks are visible to all users, but they cannot be edited.

Blockchain is becoming increasingly popular and use cases are showing up at a large variety of industries. The European Commission presented a report in April 2016 [6], supporting that blockchain technology has the potential to radically overhaul existing business models. At the same report, it was estimated that smart contracts, based on blockchain technology, could reduce infrastructure costs of banks from 13.8 to 18.4 billion euros annually by 2022. On October 2017, Bloomberg published an article which stated that Goldman Sachs and Google are among the most active blockchain investors [7]. Moreover, 10 of the largest U.S.A. banks have invested \$267 million in six blockchain companies and one consortium.

In this paper, we conduct a systematic literature review aiming to explore the business areas that blockchain technology is applied. We discuss existing or future use cases found in the literature and we analyze the impact that blockchain could have on multiple industries. We also take into consideration possible concerns that may arise from the expansion of blockchain applications to various sectors.

The remainder of the paper is structured as follows. In Sect. 2, we describe the research method, we formulate our research questions and we analyze the procedure that led us to our final set of primary studies. In Sect. 3, we report the results of the literature review and discussion follows at Sect. 4. Finally, in Sect. 5, we provide our concluding remarks.

2 Research Method

2.1 Goal and Research Questions

The goal of our research is to point out industries in which blockchain technology use cases are met. Almost ten years after it was first introduced, blockchain has expanded its use at a large variety of services beyond cryptocurrencies. This relevant experience gives us the chance to discuss issues emerged over the last years. For this, we formulated research questions, (RQ1) What are the business sectors in which blockchain applications are being used or developed? (RQ2) What are the obstacles and challenges of blockchain technology? The first question aims at discovering in the literature, applications of blockchain technology that are currently used as well as research results for other potential uses. The second question tries to showcase issues and challenges related to the expansion of blockchain application areas.

2.2 Search Process

In this paper, a systematic literature review approach was followed according to the guidelines proposed by Kitchenham [8]. In order to cover a large spectrum of relevant publications, we decided to search the following widely recognized and extensively used electronic libraries: ACM Digital Library, IEEE Xplore Digital Library, Science Direct,

and Springer Link. The keyword strings that were used are: “BLOCKCHAIN AND APPLICATIONS”, “BLOCKCHAIN AND USE CASES”. At the next step, we decided in which fields of the papers we would apply the search terms. In order to get a reasonable number of results, we searched for the keyword strings in the paper title, abstract, and keywords. We confined our search to publications written in the English language and selected as a content type only journals and conference papers, rejecting book chapters or webpages. No restrictions regarding the paper release date were used. The procedure for selecting the literature was conducted in November 2017, so it contains papers that were published or indexed up to that date. The final number of papers that we gathered after removing duplicates is 385. Then, we manually excluded papers that their title seemed irrelevant to our research, reducing the papers to 125. The above procedure was repeated by scanning the papers’ abstracts, reaching 70 papers. We read the whole text of these papers, culminating in 44 of them in order to extract information and answer our research questions (Fig. 1).

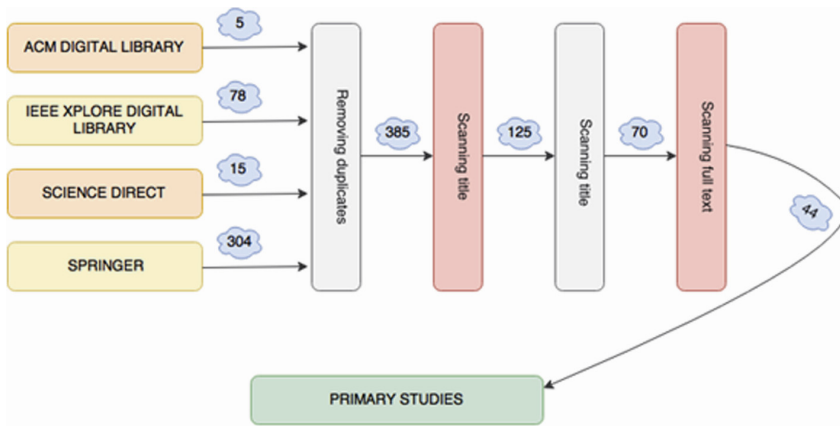


Fig. 1. Procedure for identifying primary studies

3 Results

In this section, we present the results of our literature review. The results are organized by research question and each attribute of the concept matrix is examined separately (Table 1).

Table 1. Concept matrix RQ1: business sectors in which blockchain applications are being used or developed and number of papers referring to them

	Cryptocurrencies	E-government	Healthcare	Supply chain	Energy	Banking
Total # Papers	8	14	10	7	8	5

3.1 Business Sectors in Which Blockchain Applications Are Being Used or Developed

According to our literature, we focused on six specific domains where blockchain use cases were found i.e. cryptocurrencies, e-Government, healthcare, supply chain, energy, banking, which we extensively describe below.

Cryptocurrencies. Cryptocurrencies constitute a major application area for the blockchain technology. Here, we mainly focused on the use of cryptocurrencies as a payment solution. In [2], the authors analyze the way that some of the most famous cryptocurrencies such as Bitcoin, Ethereum, Litecoin work. It is also presented a comparison among these digital currencies, regarding their coin limit, algorithm, mean block time, initial and current block rewards. In [5], the parameter of computational cost (gas) is examined and a method for reducing the gas cost while executing business processes in the Ethereum blockchain is proposed. Cryptocurrencies, could also be used as an incentive mechanism for proposing ideas in cross-functional group projects. The whole procedure is achieved via smart contract technology, which will automatically reward the group that managed to find the best idea with a predefined amount of digital coins [9]. In [19], the authors demonstrate OpenBazaar, a bitcoin-based multi-signature-protected decentralized marketplace, which enables free e-commerce transactions without any platform fees required. In [20, 32], the potential of a blockchain-assisted information distribution system for the IoT is introduced. The Internet of Things is anticipated to include sensors connected to the Internet. These devices are awaited to have access and produce a huge amount of information. Towards this, every Thing that generates an information item, may create a smart contract which will accept as an input an amount of virtual, digital coins and will output a payment receipt. The authors of [41] highlight the uses of the bitcoin blockchain protocol for payments and also use linear regression to predict attitudes towards bitcoin and the likelihood of bitcoin ownership. The implementation of a bitcoin-based community cryptocurrency is described in [42]. The suggested model includes community fund and the members may take loans that are approved by the vote of the community members.

E-government. In recent years, there is a massive expansion of e-government services to citizens, businesses and public bodies. Blockchain technology can serve as a platform capable to foster innovative applications and handle the information transactions where digitization of assets (e.g money, stocks, land properties rights) and decentralized exchange (peer to peer exchange) are involved or, could be involved. In [3, 12, 30, 31] blockchain based electronic voting systems are proposed, making votes transparent and securing that governments cannot manipulate an election because everyone is capable to read and verify the votes. The authors of [24] analyze a blockchain system that verifies the origin and genuineness of data during transmission in the e-government and public services, implemented in China. Blockchain utilizes a secure data structure that enables identifying and tracking transactions digitally and sharing the information across computer networks. In [4], the use of blockchain technology as a service support infrastructure in public sector procedures such as Digital ID management and secure document handling is discussed. The authors of [38, 40] also suggest the development of an

identity management system built on top of the Bitcoin and Ethereum blockchain respectively. The authors of [29, 33] propose innovative blockchain platforms in order to overcome the issue of tax fraud by increasing transparency. In particular, a new blockchain protocol, Pajakoin, is created as a simple, transparent and secure Value-Added Tax system [29], while in [33] a potential blockchain database is introduced, towards managing dividend flows, aiming to diminish as much as possible the double spending problem in the public taxation sector. The digitalization of the core governmental activities is likely to happen through using the blockchain platforms. A use case where academic certificates are stored in a secure way is analyzed in [26], an approach that might be very useful in the public sector. In [27], a system where private data can be shared with many organizations by the order of the user is proposed. The users of this system have full control of their data and new information is automatically updated in every organization that has access to those data. In [28], the authors suggest a novel distributed online lottery protocol that applies techniques developed for voting applications for the purpose of reducing security risks while avoiding the trusted third party. Finally, an application of blockchain is mentioned in [43] where the technology is used to track politicians' activities and serve as a transparency tool to citizens' hands.

Healthcare. Healthcare is another sector where blockchain technology could be effective. Leveraging blockchain technology, healthcare organizations could accomplish high-data volume and high-throughput transaction processing. In [3] the authors show the example of Estonia where blockchain technology is used for sharing medical records. Blockchain can be used as a way of storing and accessing medical products during the logistics process in the pharmaceutical supply chain [12]. In addition to that, blockchain can be used in sharing and managing health data securely and privately, ensuring anonymity and integrity across providers during the lifetime of a patient [13, 15, 18, 19]. The authors of [25, 39] provide a solution via blockchain to manage Electronic Medical Records in such a way that data handling of the patients becomes more secure, private and simple. More specifically, in [25], MedRec (the first and only functioning prototype where patients grant access of their personal medical information to doctors and healthcare providers), ARIA (a platform that combines radiation, medical and surgical oncology information and can assist clinicians to manage different kinds of medical data, develop oncology-specific care plans, and monitor radiation dose received by patients) and a new prototype created by the authors are presented as applications of blockchain in Healthcare. In [34], a lightweight backup and efficient recovery scheme for keys of health blockchain is suggested, whereas, in [43] a demonstration of a monitoring system is introduced where the collection of personal medical data and the notification of the patient (in case of an emergency) happen in real-time.

Energy. In [17, 20–22] blockchain technology is used in order to conduct transparent transactions in the energy market between consumers and prosumers (active consumers that both produce and consume electricity) at local energy grids consisting of renewable energy resources. In particular, the authors of [22] propose a token-based decentralized energy trading system where peers anonymously negotiate energy prices and are able to securely perform transactions. In [37] they present a local energy market scenario

with 100 residential households with artificial agents, implemented on a private blockchain, while in [3] it is highlighted that blockchain technology at local power grids allows the distribution, metering and billing of the electricity to be administered by the community itself without a reliant third-party intervention. In [11, 23] blockchain-based, intelligent, trusted measurement and monitoring of energy-related assets in a Smart Grid or a microgrid is suggested.

Supply Chain. Blockchain technology ensures identification of product provenance [14, 19] and facilitates tracking of processes [10, 36]. Furthermore, in [14], a product ownership management system is demonstrated to prevent counterfeits once the products reach the end in the supply chain. In this way, tracking of origin can be implemented after purchasing and acquiring a product. In [19] it is argued that blockchain technology provides security of supply chain. It can pinpoint the source of problematic parts and can ensure the trustworthiness between supply chain partners. Another blockchain use case in the supply chain is Everledger [26]. Everledger uses blockchain technology, which constitutes a worldwide ledger of diamonds in the luxury goods market and ensures their ownership. The authors in [43] mention that blockchain can improve the food supply chain. More specifically, they demonstrate Eattera, which is a decentralized market that connects producers with consumers by ensuring food traceability.

Banking. A number of financial institutions are currently testing transactions on blockchain platforms. Goldman Sachs, J.P Morgan, and other banking giants, have all established their own blockchain laboratories collaborating with blockchain platforms. Standard Chartered uses “Ripple”, an enterprise level blockchain platform to operate its first cross-border transactions [35, 43]. It took 10 s for the platform to complete a process that currently takes the banking system 2 days to complete. [35] also gives a thorough analysis on how blockchain can achieve asset digitization, point-to-point value transfer, thus rebuilding the financial infrastructure. This clearly increases the efficiency of clearing and settlement of financial assets after transactions. There have been estimations, that the cost of each transaction in cross-border businesses, can be extremely reduced owing to the application of blockchain. More specifically, in [10], a blockchain-based cross-border payment system is indicated, implemented in a banking blockchain platform. In addition to that, blockchain application could help banks facilitate foreign exchanges and real-time payments by gathering nodes in a blockchain, rather than having a central bank to deal with payments [3, 16]. Blockchain’s disruption in the banking sector is highlighted by IBM’s prediction that, “in 4 years, 66% of banks, will have commercial blockchain scale” [35].

Other Business Sectors. Apart from these results, blockchain can play a major role in other business sectors that are not so extensively analyzed in our literature. It can be a solution for a pay-as-you-go car insurance application. More specifically, all the data streamed from the vehicle monitoring engine are stored in the blockchain database, which guarantees that the data are tamper-proof and traceable providing a quicker and better customer experience and less operating cost as well as avoiding frauds [44]. Moreover, with the exploitation of smart contracts, blockchain can be used in farming insurance by gathering weather data, where farmers need insurance protection against

the consequences of bad weather. Smart contract could also be used, in combination with smart sensors, for home insurance with automatic reimbursements for damages [43]. Blockchain technology has also potentials for construction management. It can provide a reliable infrastructure for building information management, legal arguments and secure storage of sensor data during all life-cycle stages without using a centralized building information model avoiding the need of the trusted third party [45]. There are a lot of expectations that blockchain technology will be incredible disruptive to automotive industry as well. The future connected vehicles that will be part of IoT will need a comprehensive security architecture to protect the transferred data. In [46] the authors introduce us to an automotive security platform utilizing blockchain to tackle the implicated security and privacy challenges of future connected vehicles. Education will also be potentially benefited by the blockchain implementation. Academic credentials must be universally recognized and verifiable. Blockchain solutions in education could streamline verification procedures – thus reducing fraudulent claims of un-earned educational credits. The University of Nicosia [26] where the academic certificates of the students were stored in blockchain, is the first example moving towards this direction.

3.2 Results for RQ2 (Table 2)

Focusing on RQ2, Security, Privacy, Latency and Computational Cost are identified as the main technical challenges in the current blockchain systems found in our research.

Table 2. Concept matrix for RQ2: obstacles and challenges of blockchain technology

Papers	Privacy	Security	Latency	Computational cost
Total	6	6	3	4

Security. Even though blockchain is used in many sectors, there are still some security concerns that need to be addressed. In [2], the “Transaction Malleability” attack is described. This type of attack occurs when someone changes the unique transaction ID before the transaction is confirmed. As a result, the transaction is modified and cryptocurrencies are being transferred to the attacker’s account. The authors of [3] support that there are challenges at the blockchain’s individual nodes, whose keys may be stolen, leading to malicious transactions. In [10], it becomes clear that the most crucial issue proof of work mechanism faces is, when miners try to control more than 50% of the network’s computing power in order to prevent transactions from gaining confirmations, which is known as the 51% attack. It is summarized that further research needs to be done towards this direction, in order to find solutions that will increase blockchain’s confidence. A solution for the 51% attack is suggested at [29], where a VAT system is developed on a centralized blockchain. In [16] the authors express their concern about the fact that all nodes participating in a blockchain are connected to a P2P network but generally security leaks have been reported about P2P networks. In [14] a product ownership management system has been developed on the Ethereum blockchain but since Ethereum is still under development its security is not fully verified.

Privacy. Privacy is a main issue that is still under research. In [13], it is mentioned that even though blockchain technology can provide transparency in the clinical trial and precision medicine, this could lead to privacy concerns. The anonymity of the blockchain users' identity cannot be fully accomplished using cryptographic keys. It is also stated that in the traditional blockchain about 60% of the users' identities had been compromised via big data analytics of data sets across the Internet. Furthermore, in [16] the authors mention that financial systems, such as the banking systems, must provide high privacy in contrast to the current blockchain technology, which has a low privacy level. In [19] the authors claim that the obstacles of blockchain's deployment in the healthcare sector are also psychological, as there are data sharing concerns among medical organizations. Moreover, in some applications, privacy issues could lead to trust problems. In [31] it is claimed that in a smart grid infrastructure (energy sector) privacy is a facet of trust in the sharing economy. In addition, Bitcoin users can send digital coins to a specific address that belongs to themselves. In [49] it is mentioned that these addresses of the same person can be linked.

Latency. One of the biggest limitation all blockchain consensus protocols have in common, is that every fully participating node in the network must process every transaction. Decentralization is a core and innovative characteristic of blockchain technology but, unfortunately, drives to latency challenges. In [26], it is referred that Bitcoin blockchain theoretical transaction power is seven transactions per second. As a result, in a financial system where speed and executions in high rates are obligatory, the P2P network that blockchain provides is clearly far from being applicable [16]. The e-commerce sector is also beyond the usage of blockchain technology. The time needed for a transaction to be verified in bitcoin blockchain, which is the dominant and "traditional" blockchain protocol, is almost ten minutes something opposite to the real-time transactions that the retail businesses require. However, many platforms based on alternative blockchain protocols are developed to overcome this latency obstacle and speed up the confirmation process [50].

Computational Cost. One of the hardest issues in blockchain-based applications is the specialist hardware that is usually required, to implement a transaction via a blockchain platform [2]. Special hardware, means higher energy consumption, hence more computational costs. Due to this concentration and consolidation of mining power blockchain technologies depend upon, every potential application based on a blockchain development is imminent to further research. More precisely, in [37], a local energy market with artificial agents, implemented on a private blockchain is presented. However, the suitability of blockchain technology as mainstream ICT remains to be investigated owing to the energy consumption that incurs. Focusing in the IoT domain, authors in [44], propose the reinvention of a consensus protocol that uses tweets to encode transactions for IoT applications, to substitute the massive power that miners require towards validating the transaction. Additionally, the computational cost of blockchain technologies is highlighted in [47], where a comparison between blockchain and cloud services for software applications takes place, resulting in the fact that Ethereum's blockchain cost per process, can be a hundred times higher than on Amazon SWF.

4 Discussion

After thoroughly gathering information from our literature, in this last section we present blockchain’s hype, applications grouped by use cases, and concerns, as topics that caught our attention during conducting our review and made an impression upon us.

4.1 Blockchain’s Hype

During our research, we realized that blockchain technology is increasingly gaining in popularity and gathers huge research interest. This becomes obvious by the pie chart of the annual distribution of our primary studies (Fig. 2). About 70% of the papers that we included in our review were published in 2017.




Fig. 2. Primary studies pie chart distribution per year

According to Gartner Hype cycle of emerging technologies that was published in July 2017, blockchain technology is currently close to the borders of Peak of inflated expectations sliding slowly into the Trough of Disillusionment. In other words, the technology receives negative press for the first time. Challenges and obstacles reach the surface and implementations of the technology fail to deliver. The prediction of the firm is that the platform will be mainstream in 5 to 10 years.

Besides the fact that, introductions to cryptocurrencies have already been included in university curricula all over the world, another verification of our claim about blockchain’s disruption, is that in 2017 we had the first pedagogical attempt to teach a smart contract programming course at the University of Maryland [51]. Smart contracts can be viewed as distributed protocols executed between a number of parties. The execution of the smart contracts is automatic and it is guaranteed by the rules of the underlying cryptocurrency. The students of the course developed smart contract applications atop Ethereum blockchain using Serpent programming language.

4.2 Prospective Use Cases

Through the research we have conducted, we noticed that blockchain could undoubtedly be the foundational technology for the birth of new applications. Nevertheless, in this section, we categorize three specific areas that we believe blockchain can make a huge impact on.

Transactions/Payments. It is evident that blockchain technology is capable of creating an immutable digital ledger for transactions which can be incorporated into business processes today. Blockchain technology grants a high level of privacy by providing that transaction details are shared only amongst users involved in those transactions, thus removing the need for a central authority to administer them. As depicted above, in papers [9, 19] we highlight some specific applications based on a blockchain platform, where cryptocurrencies are used as a payment solution. Moreover, in the E-government sector, increasing transparency in the transactions, which is a blockchain's platform advantage, could help overcome issues of tax fraud, as indicated in [29, 33]. Finally, considering the ability of clearing and the agreement of the financial assets after transactions that blockchain platforms provide, we point out the foundational disruption of blockchain technology in the banking industry, illustrating certain banking blockchain platforms, found in papers [10, 35, 43].

Data Storage. Blockchain technology and its special features could bring big advantages to data storage systems, as it will provide cheaper, faster, more secure and decentralized storage than the existing cloud storage platforms. Decentralized storage works by distributing the data across a network of nodes, in a similar way to the distributed ledger technology characteristic of blockchain. Blockchain applications that already encompass decentralized storage are Storj, Madsafe and IPFS [43]. The use of blockchain as a database applies to different economic sectors. E-government, healthcare, banking and supply chain sectors try to implement various blockchain projects to optimize their operational procedures. MedRec, ARIA [3], University of Nicosia certificate storage, Bitnation, E-resident and Everledger [26] are only some of those projects that their usage could bring enormous profits. On the contrary, having a great potential is not the same as having great success. There are still some insurmountable obstacles that prevent the adoption of the technology with its current structure. New protocols and platforms are created every single day to correct the previous ones. There is no doubt that blockchain in the near future will make interaction between people and organizations faster and cheaper.

ID Management. Blockchain technology can become a powerful tool for identity management. As we are continuously being asked to share personal information to access places or information or to do business with other companies, we are at risk for identity theft. Blockchain constitutes the underlying technology for identity management through decentralized networks. As shown in [24], in e-government applications, identity management with blockchain can provide each citizen with a verifiable digital immutable identity, simplifying processes and improving the speed and authority in government approval. Furthermore, it is found that blockchain can facilitate patients' health identity management by giving pharmacists and doctors access to patients' electronic medical records [15, 18]. This would allow care providers, pharmacists and patients to track dosages, receive automatic alerts for missed or incorrect dosages, monitor possible adverse drug interactions and even help prevent addiction.

4.3 Usage Concerns

Trying to elaborate on blockchain's issues and challenges, we mainly focused on the vulnerabilities that came up with blockchain's expansion to new areas of implementation. Each application have extra requirements the existing blockchains fail to meet. Research recognizes those inconsistencies and suggests customized blockchain solutions. A great example of how current blockchains are inadequately effective is the time that it takes for a transaction to be completed. When bitcoin first started, money transactions that would probably take days to complete, were carried out in a few minutes and that was a revolution. But, as we see in [16, 50] stock markets or e-commerce applications have transactions that have to be completed almost instantly.

Applications in different business areas also demand different levels of privacy. Current public blockchains have low privacy but future use cases in banking systems [16] or in medical records in the healthcare sector [19] require high privacy. Security is another domain where increasing concerns are expressed. The computers' increasing computational power along with the rise of mining pools (groups of people mining together as a single unit) could result in an attack on the blockchain if somehow, someone was able to control 51% of the mining power. This scenario is discussed in [10, 16, 29] and even though it does not seem plausible, a potential attack would be fatal for the blockchain's reputation. For example, a 51% attack on the bitcoin could significantly devalue it. Blockchain's increasing popularity has as a result the boost of miners' number. More and more people are using their computational power to get a cryptocurrency reward. Since only the miner who solves the transaction encryption first gets the reward, all the others just waste resources [37, 47].

The blockchain industry is currently receiving huge attention from everyday startups or tech-people. However, given the fact that the idea of implementing a blockchain is still in the introductory stage of its life cycle and requires a set of skills and knowledge which are not feasible easily, here we present some roadblocks that make its mass adoption truly ambitious, at least at this stage.

As it is known, by disintermediate financial institutions, multiple parties are able to conduct transactions easily without paying a commission. Technically speaking, moving cash to a blockchain infrastructure could lead to a significant increase of the overall transaction cost or, trading on a blockchain system would also be slower than traders would tolerate and mistakes may be inevitable, potentially bringing huge losses. In addition to that, due to its genetically distributed peer-to-peer nature, blockchain transaction can only be completed when all parties update their corresponding ledgers, a process that might take hours. This transaction delay may be a deal-killer. The difficulty of a mass implementation of blockchain technology is visible if we consider that the commitment of blockchain in large part depends upon enough parties using the same implementation of the technology, requiring a universal adoption. Blockchain technology has not yet imbued into many real use cases and besides the technical part, we would like to highlight, that maybe the biggest obstacle of a mass implementation is the education and knowledge that is required even for someone with strong tech-background, in order to fully understand the benefits of this new area of technology. Consider that even the terminology of blockchain is too complex and creates itself some obstacles

and doubts, even to CEO's, let alone to everyday consumers. It takes time for a new technology, especially to something as foundational as blockchain, to incorporate itself into the fabric of modern society. In fact, "the more likely blockchain is to disrupt the global financial system, the less likely is to succeed" [48].

Another challenge that has to be tackled about blockchain is establishing standardization and regulatory framework. In April 2016, European Union in one of its reports points out that "The future of blockchain requires the development of a common language with specific rules for interaction, which will be achieved through standardization processes." [8]. In September 2016 ISO accepted Australia's proposal to manage the Secretariat of ISO/TC 307 for new international standards on blockchain and in March 2017 the Roadmap Report was released [52, 53]. Currently, 29 participating member-states and 13 observing member-states are developing 4 standards that are at proposal or preparatory stage.

For the time being, the existing legal framework applies to activities related with the blockchain technology. The regulators are monitoring blockchain-based activities, acquiring knowledge in order to make the law keep up with the evolution of the technology. Regulators are gradually starting to understand the blockchain use cases, while at the same time the innovators are trying to find the regulatory principles that apply to their activity. Current blockchain applications are considered rather immature to guide the legislation towards a specific direction. The development of the existing applications and the creation of new ones are expected to showcase the legal gaps that need to be regulated. In the meantime, existing legal principles can sufficiently face blockchain-related criminal activity whereas experience gained over time will provide guidance and dictate the need for new regulations that cover all the legal scenarios that may arise by the use of blockchain technology.

5 Conclusion

Blockchain technology is challenging the status quo in several areas in a radical way providing a decentralized database of any transaction involving transfer of value i.e. (money, goods, property, assets or even votes). This generic nature is what makes blockchain technology attractive to many business areas today. However, we conclude our review highlighting the risks, effects as well as the unintended consequences of blockchain technology on established markets. The disruption of blockchain technology in business sectors is increasing in pace. Therefore, we believe that further critical research is needed to exploit its capabilities and understand the limitations when applied in a large scale.

References

1. Nakamoto, S.: Bitcoin: a peer-to-peer electronic cash system (2008)
2. Suhaliana bt Abd Halim, N., Rahman, M.A., Azad, S., Kabir, M.N.: Blockchain security hole: issues and solutions. In: Saeed, F., Gazem, N., Patnaik, S., Saed Balaid, A.S., Mohammed, F. (eds.) *IRICT 2017. LNDECT*, vol. 5, pp. 739–746. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-59427-9_76
3. Bhardwaj, S., Kaushik, M.: Blockchain—technology to drive the future. In: Satapathy, S.C., Bhateja, V., Das, S. (eds.) *Smart Computing and Informatics. SIST*, vol. 78, pp. 263–271. Springer, Singapore (2018). https://doi.org/10.1007/978-3-319-59427-9_76
4. Ølnes, S., Jansen, A.: Blockchain technology as a support infrastructure in e-government. In: Janssen, M., Axelsson, K., Glassey, O., Klievink, B., Krimmer, R., Lindgren, I., Parycek, P., Scholl, Hans J., Trutnev, D. (eds.) *EGOV 2017. LNCS*, vol. 10428, pp. 215–227. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-64677-0_18
5. García-Bañuelos, L., Ponomarev, A., Dumas, M., Weber, I.: Optimized execution of business processes on blockchain. In: Carmona, J., Engels, G., Kumar, A. (eds.) *BPM 2017. LNCS*, vol. 10445, pp. 130–146. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-65000-5_8
6. Probst, L., Frideres, L., Cambier, B., Martinez-Diaz, C.: PwC Luxemburg: Blockchain Applications and Services, Directorate-General for Internal Market, Industry, Entrepreneurship and SMEs, European Union (2016)
7. Bloomberg Article. <https://www.bloomberg.com/news/articles/2017-10-17/goldman-google-make-list-of-most-active-blockchain-investors>. Accessed 21 Dec 2018
8. Keele, S.: Guidelines for performing systematic literature reviews in software engineering. Technical report, Ver. 2.3 EBSE Technical report. EBSE (2007)
9. O’Leary, K., O’Reilly, P., Feller, J., Gleasure, R., Li, S., Cristoforo, J.: Exploring the application of blockchain technology to combat the effects of social loafing in cross functional group projects. In: *Proceedings of the 13th International Symposium on Open Collaboration - OpenSym 2017* (2017). <https://doi.org/10.1145/3125433.3125464>
10. Wu, T., Liang, X.: Exploration and practice of inter-bank application based on blockchain. In: *2017 12th International Conference on Computer Science and Education (ICCSE)* (2017)
11. Wu, L., Meng, K., Xu, S., Li, S., Ding, M., Suo, Y.: Democratic centralism: a hybrid blockchain architecture and its applications in energy internet. In: *2017 IEEE International Conference on Energy Internet (ICEI)* (2017). <https://doi.org/10.1109/ICEI.2017.38>
12. Bocek, T., Rodrigues, B., Strasser, T., Stiller, B.: Blockchains everywhere - a use-case of blockchains in the pharma supply-chain. In: *2017 IFIP/IEEE Symposium on Integrated Network and Service Management (IM)* (2017). <https://doi.org/10.23919/INM.2017.7987376>
13. Shae, Z., Tsai, J.: On the design of a blockchain platform for clinical trial and precision medicine. In: *2017 IEEE 37th International Conference on Distributed Computing Systems (ICDCS)* (2017). <https://doi.org/10.1109/ICDCS.2017.61>
14. Toyoda, K., Mathiopoulos, P., Sasase, I., Ohtsuki, T.: A novel blockchain-based Product Ownership Management System (POMS) for anti-counterfeits in the post supply chain. *IEEE Access* **5**, 17465–17477 (2017). <https://doi.org/10.1109/ACCESS.2017.2720760>
15. Zhang, J., Xue, N., Huang, X.: A secure system for pervasive social network-based healthcare. *IEEE Access* **4**, 9239–9250 (2016)
16. Tsai, W., Blower, R., Zhu, Y., Yu, L.: A system view of financial blockchains. In: *2016 IEEE Symposium on Service-Oriented System Engineering (SOSE)* (2016)

17. Munsing, E., Mather, J., Moura, S.: Blockchains for decentralized optimization of energy resources in microgrid networks. In: 2017 IEEE Conference on Control Technology and Applications (CCTA) (2017). <https://doi.org/10.1109/CCTA.2017.8062773>
18. Mettler, M.: Blockchain technology in healthcare: the revolution starts here. In: 2016 IEEE 18th International Conference on e-Health Networking, Applications and Services (2016)
19. Kshetri, N.: Blockchain's roles in strengthening cybersecurity and protecting privacy. *Telecommun. Policy* **41**, 1027–1038 (2017). <https://doi.org/10.1016/j.telpol.2017.09.003>
20. Sikorski, J., Haughton, J., Kraft, M.: Blockchain technology in the chemical industry: machine-to-machine electricity market. *Appl. Energy* **195**, 234–246 (2017)
21. Castellanos, J., Coll-Mayor, D., Notholt, J.: Cryptocurrency as guarantees of origin: simulating a green certificate market with the Ethereum Blockchain. In: 2017 IEEE International Conference on Smart Energy Grid Engineering (SEGE) (2017)
22. Aitzhan, N.Z., Svetinovic, D.: Security and privacy in decentralized energy trading through multi-signatures, blockchain and anonymous messaging streams. *IEEE Trans. Dependable Secure Comput.* **1** (2016). <https://doi.org/10.1109/TDSC.2016.2616861>
23. Imbault, F., Swiatek, M., de Beaufort, R., Plana, R.: The green blockchain: managing decentralized energy production and consumption. In: 2017 IEEE International Conference on Environment and Electrical Engineering and 2017 IEEE Industrial and Commercial Power Systems Europe (EEEIC/I&CPS Europe) (2017)
24. Hou, H.: The application of blockchain technology in e-government in China. In: 2017 26th International Conference on Computer Communication and Networks (ICCCN) (2017)
25. Dubovitskaya, A., Xu, Z., Ryu, S., Schumacher, M., Wang, F.: How blockchain could empower ehealth: an application for radiation oncology. In: Begoli, E., Wang, F., Luo, G. (eds.) DMAH 2017. LNCS, vol. 10494, pp. 3–6. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-67186-4_1
26. Ølnes, S.: Beyond bitcoin enabling smart government using blockchain technology. In: Scholl, H.J., Glassey, O., Janssen, M., Klievink, B., Lindgren, I., Parycek, P., Tambouris, E., Wimmer, Maria A., Janowski, T., Sá Soares, D. (eds.) EGOVIS 2016. LNCS, vol. 9820, pp. 253–264. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-44421-5_20
27. Alboaie, S., Cosovan, D.: Private data system enabling self-sovereign storage managed by executable choreographies. In: Chen, L.Y., Reiser, H.P. (eds.) DAIS 2017. LNCS, vol. 10320, pp. 83–98. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-59665-5_6
28. Grumbach, S., Riemann, R.: Distributed random process for a large-scale peer-to-peer lottery. In: Chen, L.Y., Reiser, H.P. (eds.) DAIS 2017. LNCS, vol. 10320, pp. 34–48. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-59665-5_3
29. Wijaya, D.A., Liu, J.K., Suwarsono, D.A., Zhang, P.: A new blockchain-based value-added tax system. In: Okamoto, T., Yu, Y., Au, M.H., Li, Y. (eds.) ProvSec 2017. LNCS, vol. 10592, pp. 471–486. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-68637-0_28
30. Zhao, Z., Chan, T.-H.H.: How to vote privately using bitcoin. In: Qing, S., Okamoto, E., Kim, K., Liu, D. (eds.) ICICS 2015. LNCS, vol. 9543, pp. 82–96. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-29814-6_8
31. Sun, J., Yan, J., Zhang, K.: Blockchain-based sharing services: what blockchain technology can contribute to smart cities. *Finan. Innov.* **2** (2016)
32. Polyzos, G., Fotiou, N.: Blockchain-assisted information distribution for the Internet of Things. In: 2017 IEEE International Conference on Information Reuse and Integration (IRI)
33. Hyvärinen, H., Risius, M., Friis, G.: A blockchain-based approach towards overcoming financial fraud in public sector services. *Bus. Inform. Syst. Eng.* **59**, 441–456 (2017). <https://doi.org/10.1007/s12599-017-0502-4>

34. Zhao, H., Zhang, Y., Peng, Y., Xu, R.: Lightweight backup and efficient recovery scheme for health blockchain keys. In: 2017 IEEE 13th International Symposium on Autonomous Decentralized System (ISADS) (2017). <https://doi.org/10.1109/ISADS.2017.22>
35. Guo, Y., Liang, C.: Blockchain application and outlook in the banking industry. *Finan. Innov.* **2** (2016). <https://doi.org/10.1186/s40854-016-0034-9>
36. Zhao, J., Fan, S., Yan, J.: Overview of business innovations and research opportunities in blockchain and introduction to the special issue. *Finan. Innov.* **2** (2016)
37. Mengelkamp, E., Notheisen, B., Beer, C., Dauer, D., Weinhardt, C.: A blockchain-based smart grid: towards sustainable local energy markets. *Comput. Sci. Res. Dev.* **33**, 207–214 (2017). <https://doi.org/10.1007/s00450-017-0360-9>
38. Augot, D., Chabanne, H., Chenevier, T., George, W., Lambert, L.: A user-centric system for verified identities on the bitcoin blockchain. In: Garcia-Alfaro, J., Navarro-Arribas, G., Hartenstein, H., Herrera-Joancomartí, J. (eds.) ESORICS/DPM/CBT -2017. LNCS, vol. 10436, pp. 390–407. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-67816-0_22
39. Liu, P.T.S.: Medical record system using blockchain, big data and tokenization. In: Lam, K.-Y., Chi, C.-H., Qing, S. (eds.) ICICS 2016. LNCS, vol. 9977, pp. 254–261. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-50011-9_20
40. Azouvi, S., Al-Bassam, M., Meiklejohn, S.: Who Am I? secure identity registration on distributed ledgers. In: Garcia-Alfaro, J., Navarro-Arribas, G., Hartenstein, H., Herrera-Joancomartí, J. (eds.) ESORICS/DPM/CBT-2017. LNCS, vol. 10436, pp. 373–389. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-67816-0_21
41. Bashir, M., Strickland, B., Bohr, J.: What motivates people to use bitcoin? In: Spiro, E., Ahn, Y.-Y. (eds.) SocInfo 2016. LNCS, vol. 10047, pp. 347–367. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-47874-6_25
42. Vandervort, D., Gaucas, D., Jacques, R.S.: Issues in designing a bitcoin-like community currency. In: Brenner, M., Christin, N., Johnson, B., Rohloff, K. (eds.) FC 2015. LNCS, vol. 8976, pp. 78–91. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-48051-9_6
43. Lamberti, F., Gatteschi, V., Demartini, C., Pranteda, C., Santamaria, V.: Blockchain or not blockchain, that is the question of the insurance and other sectors. *IT Professional* **1** (2017)
44. Vo, H., Mehedy, L., Mohania, M., Abebe, E.: Blockchain-based data management and analytics for micro-insurance applications. In: Proceedings of the 2017 ACM on Conference on Information and Knowledge Management - CIKM 2017 (2017)
45. Turk, Ž., Klinc, R.: Potentials of blockchain technology for construction management. *Procedia Eng.* **196**, 638–645 (2017). <https://doi.org/10.1016/j.proeng.2017.08.052>
46. Steger, M., Dorri, A., Kanhere, S., Römer, K., Jurdak, R., Karner, M.: Secure Wireless Automotive Software Updates Using Blockchains: a proof of concept. *Adv. Microsyst. Autom. Appl.* **2017**, 137–149 (2017)
47. Rimba, P., Tran, A., Weber, I., Staples, M., Ponomarev, A., Xu, X.: Comparing blockchain and cloud services for business process execution. In: 2017 IEEE International Conference on Software Architecture (ICSA) (2017). <https://doi.org/10.1109/ICSA.2017.44>
48. Forbes article. <https://www.forbes.com/sites/quora/2017/09/21/whats-holding-blockchain-back-from-large-scale-adoption/#267d559d2309>
49. Conoscenti, M., Vetrò, A., De Martin, J.C.: Blockchain for the Internet of Things: a systematic literature review, pp. 1–6 (2016)
50. Xu, Y., Li, Q., Min, X., Cui, L., Xiao, Z., Kong, L.: E-commerce blockchain consensus mechanism for supporting high-throughput and real-time transaction. In: Wang, S., Zhou, A. (eds.) CollaborateCom 2016. LNICST, vol. 201, pp. 490–496. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-59288-6_46

51. Delmolino, K., Arnett, M., Kosba, A., Miller, A., Shi, E.: Step by step towards creating a safe smart contract: lessons and insights from a cryptocurrency lab. In: Clark, J., Meiklejohn, S., Ryan, P.Y.A., Wallach, D., Brenner, M., Rohloff, K. (eds.) FC 2016. LNCS, vol. 9604, pp. 79–94. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53357-4_6
52. ISO. <https://www.iso.org/committee/6266604.html>
53. Roadmap for blockchain standards – Report. http://www.standards.org.au/OurOrganisation/News/Documents/Roadmap_for_Blockchain_Standards_report.pdf



ICT-Based Support for the Collaboration of Formal and Informal Caregivers – A User-Centered Design Study

Madeleine Renyi^{1,2(✉)}, Frank Teuteberg², and Christophe Kunze¹

¹ Technology and Inclusion Lab, Furtwangen University, Furtwangen, Germany
{madeleine.renyi, christophe.kunze}@hs-furtwangen.de

² School of Business Administration and Economics,
Department of Accounting and Information Systems,
Osnabrück University, Osnabrück, Germany
frank.teuteberg@uni-osnabrueck.de

Abstract. Given the demographic change and the resulting need for comprehensive care strategies, collaborative care plays an important role in ensuring care for everyone in need in the future. Despite the fact that collaboration software tools can significantly relieve caregivers by improving various work processes, their use in informal care networks has not yet become common practice. For this purpose, a mixed method user-centered design study including literature review, market research, system analysis, case studies as well as workshops and interviews were conducted. On the basis of this, the needs and requirements for mobile collaboration support were identified. The design of the requested application must strike the balance between the request for simple functionality and the diversity of information needed in order to offer an added value for all participants.

Keywords: E-health · Computer supported collaborative work · Informal care
Care-mix · Mobile application

1 Introduction

The e-health business sector is steadily driving innovation in the healthcare market. A special field of application of e-health is the support of long-term care. In principle, long-term care is divided into four sectors [1]: (1) family and informal care sector, (2) state or public sector, (3) voluntary and non-governmental-organization (NGO) sector and (4) care market or private sector, whereby the respective shares depend on country, family situation and health status of the person in need. As elderly care basically represents an individual mixture of different care services, in the remainder of this contribution, we refer to elderly care as care-mix. Bäuerle and Scherzer [2] define care-mix as the composition of various support services and daily activities at an individual level, whereby professional and voluntary services as well as help from relatives are combined and networked.

There are several reasons illustrating that care-mix arrangements are important for ensuring nursing care, for instance, the increasing burden on caring relatives

(compatibility of work and long-term care) and the expected increasing lack of skilled employees in professional care [3]. In addition to nursing activities, care-mix also includes services such as driving, meals on wheels or neighborhood assistance. In this context, approaches like organized neighborhoods and caring communities are becoming more and more the center of attention.

The future role of technical support systems in care-mix settings has not been investigated in depth so far. Görres et al. [3] argue that the potential of computer supported organized work is currently far from being fully exploited. The goal of the technology deployment must be to coordinate and control the individual offers in order to optimize care processes. As pointed out by Pinelle and Gutwin [4], present care-mix configurations are usually weakly structured and loosely coupled work processes in which the involved employees have only limited access to information from other organizations. Furthermore, the current low level of networking between the various actors impedes the introduction of cross-organizational processes and information exchange systems. The difficulty of reaching a “critical mass” of users, which is necessary for a meaningful system use, as well as unequally distributed added values are described as typical barriers [4].

The underlying research questions (RQ) of this contribution therefore are:

- RQ1: What are the user expectations and requirements towards a support system for collaboration in care-mix settings?
- RQ2: What are the specific components of a mobile system supporting collaboration in care-mix settings?

This article presents a comprehensive needs assessment and system design approaches for a collaboration software to support care-mix teams. The article is structured as follows: The research method is set out in Sect. 2. Section 3 provides an overview over related work. The results and possible answers to the research questions are presented in Sect. 4. Section 5 shows a possible implementation strategy. The contribution is complemented by the discussion of the implications for science and practice, the resulting conclusions as well as the perspectives for future work in this project (Sect. 6).

2 Study Design and Methods

To obtain a deep understanding of the “organizational culture” [5], an iterative mixed method approach was chosen. Methods of observation and understanding (literature reviews, interviews, market research, system analysis and case studies) were combined with a participatory design process [6]. The aim was to derive the requirements for digital support of care-mix settings and to evaluate these. For the purpose of identifying appropriate needs and requirements, participants of social systems were integrated by means of different workshops. The findings subsequently served as a basis for concept creation, concept evaluation and concept review.

This iterative process was passed through in three different regions in the German state of Baden-Württemberg: the rural municipality Fischerbach and the two urban

regions Kirchheim unter Teck and Wernau. In Fischerbach and Wernau there are well organized and highly frequented neighborhood associations. Although in Kirchheim unter Teck there are no such structures, the local district Rauner participates in a funded project for the construction of district networking structures. In total 12 interviews and 10 workshops were conducted in these regions.

Interviews: The results of the literature analysis were used to develop guidelines that enable a standardized interview conduction, which again ensures verifiable and usable interview results. The interviews were conducted with two caring relatives, two civic volunteers, one professional nurse, two coordinators of organized neighborhood assistance, three coordinators of care services, one coordinator of a welfare association and one district manager. The interviews lasted between 30 and 70 min. All interviews were recorded, transcribed and analyzed using the computer software MAXQDA (<http://www.maxqda.de/>). The data-analyzation method is based on the concept of Mayring [7].

Case Studies: To better understand the user requirements towards communication systems in nursing networks, two “technology probe” case studies were performed based on the approach of Hutchinson et al. [8]. His basic assumption is that social practices change with the usage of technology. Therefore, pure observations of the current situation are not expedient to comprehensively assess the needs. However, by using technological artifacts, the resultant changes can be recorded and incorporated into the further concept formulation. Through the subsequent market investigation, the SimpliCare app (<http://www.simplicare.net>) was chosen as suitable for the envisaged case studies with outpatient care networks. SimpliCare is drawn up in German-language and easily available via the Google Play Store. Two “technology probe” case studies were prepared and carried out in the rural municipality Fischerbach in Germany (test duration Nov. 2015 through Feb. 2016). Network 1 span around a 93-year-old lady suffering from dementia at an early stage. The care network consisted of her daughter and son in law as well as a volunteer of the neighborhood assistance association. The second network consists of an 84-year-old lady suffering from dementia, her sister and nephew, an employee of the local welfare center and a volunteer of the neighborhood assistance association.

Workshops: Using a scenario-based design approach according to Rosson and Carrol [9], the results of the literature review and the evaluation of the interviews as well as the experiences of the experimental case studies were transferred in problem scenarios which were presented and discussed in stakeholder workshops. These small group discussions provided more information on experiences, views and needs. In order to involve future users in the concept development for the later prototype, concept creation workshops were carried out. The ideas were transferred into interactive mock-ups and evaluated in further workshops.

3 Related Work

The capabilities of supporting home-care settings through digital and connecting technologies are discussed in most developed countries. The findings from literature relate to

the fields documentation, (activity and social) awareness and collaboration, whereas the presented systems include a wide range of recent technologies like mobile, ubiquitous and ambient computing. Fitzpatrick and Ellingsen [10] provide insights of 25 years (1987–2012) of computer supported collaborative work research in healthcare.

Functionalities: An approach to support the care-mix in outpatient care is for example presented by Pinelle and Gutwin [4]. The groupware system Mohoc is used to exchange information on activities, news and group discussions within cross-organizational, loosely coupled home care teams. The collaboration in outpatient care teams is rather infrequent and depends on the respective situation of the care recipient. The features supporting autonomous work activities were used most frequently in the Mohoc field test and were evaluated as generally positive and supportive for loosely coupled tasks. While the Mohoc system is mainly designed to support professional care actors, the Danish CareCoor [11] system is laid out to equally support informal (family members) and professional actors (home care workers) of one care setting. CareCoor is a tablet- and web-based tool that supports the coordination of tasks and appointments as well as the exchange of messages in care networks. By means of a shared patient calendar, all involved obtain an overview of upcoming, already performed or due activities. Findings of the first field trials showed that a responsive smartphone version would supplement the platform, which was in general rated positively by all actors. Due to the fact that many informal caregivers are also elderly and need help to some degree as well, Moser et al. [12] conducted an extensive requirements analysis for informal caregivers in order to design a platform on which informal caregivers can connect to other people to offer or request for help. In several interviews, the subject cross-generational contact and the need for an intermediary, as contact person for all kinds of issues, were stressed.

All contributions equally stress the importance of preparing and conducting field studies in the respective environment, as this is the only way to reveal unforeseen processes and implementation obstacles.

Design: Due to the fact that health care systems vary, home care is managed differently from country to country. Therefore, research findings of one country may not hold true for another. And yet, general design guidelines for informal care collaboration may be useful (cf. [13, 14]).

Business Models: Although in literature there is no detailed information about suitable business models, the market research revealed various approaches.

Financing by end users: In this approach, the primary caregivers bear the costs. Each access to the application is charged. While receiving the full range of functions, the user can create as many networks as needed. This business model is for example used for the SimpliCare app, which was used in the case studies. However, download statistics of the Google Play store indicate that this model may not be of great success.¹

Financing by care provider: In order to differentiate from competitors and increase customer loyalty, home care providers offer their own apps to improve the networking

¹ Installation 10-50, last accessed: 13.12.2017, <https://play.google.com/store/apps/details?id=com.begsolutions.simplicare.full>.

with family caregivers and care recipients and thus enhance transparency. Examples for this business model are Mavencare (<https://mavencare.com/connected-home-care>) and honor (<https://www.joinhonor.com/apps>).

Visionaries like Buurtzorg (Netherlands, <http://www.buurtzorgnederland.com/>) see nursing and caring from a community perspective. This new organizational model connects autonomous, self-guided teams with the neighborhood via cross-generational platforms in order to ensure needs-oriented care structures. This concept has revolutionized the care market in the Netherlands and has already inspired other countries (e.g. Belvita, Switzerland, <https://belvita.ch/belvita-idee/>).

Mixed financing: An example for this model is the Jointly-App (<https://jointlyapp.com/>). Payment is made per care circle for which all associated accounts are free. The fee for one circle (\$ 2.99) can be individually borne by care providers, companies (compatibility of work and family), health insurances or informal caregivers. The present download statistics indicate that there is a certain interest in this business model.²

Beyond previous findings in the literature, this article aims to contribute to the understanding of collaboration needs in care-mix settings in Germany. According to the authors, it is indispensable to follow a holistic, mixed-method approach in order to involve all participants in outpatient care networks. Further, the design of a sustainable, functioning business model imperatively requires the involvement of the public and private sector. Thereby, a special focus lies on the location- and device-independent mobile support of the target group.

4 Results

4.1 Observation and Understanding of Collaboration in Care-Mix Teams

The findings gained from our interviews and workshops are in line with those from literature (c.f. related work section).

Technology and Usability: The technical affinity among the participants varies considerably, which is dependent on age and (professional) experience of the respective user. By and large, the affinity for technology within the target group is rather low and ranges from daily (work and private life) to no use at all. And even though all respondents stated to possess a smartphone, they mainly use it to make phone calls. In fact, phone calls and e-mails constitute the preferred communication channels for professional caregivers. Semi-professionals and informal caregivers favor instant messaging services for fast communication.

The preliminary assumption that younger and technologically more experienced actors would be especially open and interested in the use of such platforms could not be confirmed. Furthermore, the case studies showed that particularly less technically experienced caregivers still have difficulties in handling mobile communication systems while technically versed caregivers almost intuitively operate standard tools to improve

² Installation 1.000-5.000, last accessed: 13.12.2017, <https://play.google.com/store/apps/details?id=org.carersuk.jointlyapp>.

communication and cooperation. The generally very appealing design as well as the quite high usability of common tools set the standard for the care-specific software solutions. If the specific tools do not meet the expectations, it must be reckoned that these will be rejected by technically affine users. From the second case study it emerged that the users perceived the different message types (tasks, appointments, text messages) as rather confusing.

Responsibilities: When family members live in the immediate vicinity, they usually take over the coordination of the individual care case as caregiving relative. Otherwise, the role of the coordinator is taken by an official supervisor. The primary responsible person in a care-mix setting has to keep the overview, is the contact person in case of questions or problems, manages the care-mix and leads the network of actors.

According to the participants, networking the different actors would enhance the care process. Usually, the respective actors know nothing about each other or who is responsible for what. Here, a consistent networking would ameliorate the coordination of appointments, foster social interaction and strengthen the capacity for emergency prevention and management.

Appointments: As is stated in literature [15] and confirmed in our interviews and workshops, calendars constitute an essential coordination tool in home care. In almost any care setting, there is a paper calendar as central element used for information purposes, coordination of appointments and scheduling the different services.

Nevertheless, the actors complain about conflicting schedules due to a missing topicality of the calendar as well as the difficulties with the coordination of shared appointments.

Communication: Communication usually takes place directly between the primary responsible and the other actors in the care-mix, whereby phone and e-mail are the main communication tools. In case the target person is not instantly available, repeated contact attempts by phone can be time-consuming. E-mail is not applicable for time-sensitive, urgent topics as response times are usually longer. Civic volunteers and organized neighborhood helpers also favor instant messaging tools for fast communication. Another often used means of communication are paper notes, placed at strategic places in the care recipient's household. However, this communication method lacks any read confirmation, which is criticized by several participants. As there is virtually no communication between the others in the care network, an overview of all involved would be needed.

In general, there are two main causes requiring information exchanges: (1) changes in the usual process sequences of the care-mix arrangement (e.g. deterioration of the health status of the care recipient) and (2) the continuous flow of information. A regular information update on the individual clients is desirable for all participants. A not notified absence of the patient poses a serious problem for all actors as it cannot be excluded that there is an emergency.

“Tacit” Case Knowledge: Care-mix settings with professional nursing care usually benefit from a patient record with a report section in the flat of the care recipient. This

record simplifies communication as all client-related information (changes in the health status, peculiarities in the care, etc.) can be entered and thus made available for the other actors. While organized neighborhood helpers receive information by their managers if such a file does not exist, civic volunteers without a connection to a welfare center or organized neighborhood association only receive information via direct communication with the primarily responsible.

Although all actors consider it very important to be kept informed, case conferences between informal, semi-professional and professional actors are uncommon.

4.2 Participatory Design Process - Concept Creation of a Care-Mix Collaboration Tool

Also Ganoë et al. [16] emphasize the importance of providing awareness of the overall situation, including dependencies and shared knowledge, within a collaborative group. Therefore, the designed application aims at providing a functioning information sharing system that does not only increase the actors’ awareness of the on-going activities, but at the same time enables mutual support. An aggregated list of requirements is presented in Table 1. Due to the diversity of the actors in care-mix settings, the design and arrangement of the functionalities must be as simple as possible, while a large amount

Table 1. Aggregated requirements for a collaboration tool for care-mix teams

Functional requirements	lr	i	cs	ws
Personal account for care recipient	[12]			•
Create, edit, delete own user profile	[12]			•
Provide an overview over all networks	[16]			•
Create, edit, withdraw, delete network (role dependent)	[13]	•	2	•
Case-related knowledge (including medication, etc.)	[4, 16]	•		•
Case-related directory (containing field of work, preferred communication channel, availability times, etc.)	[4, 12, 13]	•		•
Create, edit, delete appointments and tasks	[4, 11, 12, 14–16]	•		•
Create, edit, delete pin board entry (including comment, like, tag, filter entries)	[13, 16]		2	•
Send (case-related) (group) messages	[4, 11–13, 16, 17]			•
Push notifications for all news/actions			1, 2	•
Case-related and personal notes			2	•
Online and offline functionality			•	•
Non-functional requirements	lr	i	cs	ws
Mobile application	[4, 11, 17]		2	•
Networking of all actors	[16]	•		•
Performance			2	•
Usability		•	1, 2	•
Confidentiality, integrity and data security	[12, 15, 17]			•

^aNote: literature review (lr), interviews (i), case study (cs), workshop (ws)

of diverse information is to be organized. In this, it seems appropriate to arrange the functionalities network-specific. By entering the application the user gets an overview of all care networks in which he is involved.

Instant Messenger for Direct Communication: To avoid time consuming phone calls, an instant messenger allows for asynchronous communication. As some actors already eagerly use instant messaging services to communicate within their care-mix networks, it is self-evident that design and usability of the newly designed messenger must mandatorily meet the standard of the existing services to be accepted by the future users. Group discussions as well as peer-to-peer communication must be possible. Chats are assigned to a specific network, for a better clarity and allocation.

Pin Board for Informal Communication: The pin board is meant to replace the previous communication via paper notes. In this way information “not worth a call” can be conveyed in the respective care-mix networks. For a fast overview of all entries in all networks a separate collected view is needed. Via one click a private chat with the author of the entry can be opened. Moreover, a comment and a like function encourage social interaction. Tasks from the pin board can be marked as “done”. Appointments can be created from pin board entries.

Directories for Clear Responsibilities and Awareness of Care Actors: Unclear responsibilities and missing knowledge are avoided through network specific directories. The address book contains all contacts that are important for the person in need of care.

Information Page for Case Knowledge: To gain a basic knowledge of the care recipient, an information page is provided that contains e.g. the housing situation or preferences of the care recipient. This page does not replace care documentation of a care service provider, and only contains non-sensitive personal data to enhance the collaboration in the care-mix team.

Calendar for the Coordination of Appointments: The network specific calendar can be used to organize appointments and tasks concerning the care of the central network person, which prevents conflicting schedules. By means of an export function, appointments can be transferred to the personal calendars.

5 Architecture for a Mobile Support Application

All respondents explicitly stress that the use of smartphones would be desirable. Notwithstanding the fact that most of the respondents use android-devices, the aspired collaboration tool must equally be suitable for all involved actors. Hence, the application is to be programmed device-independent (cf. client tier in Fig. 1). In this way, no one is obliged to acquire new and expensive devices, which is what some respondents worry about and may lead to refusal of the tool.

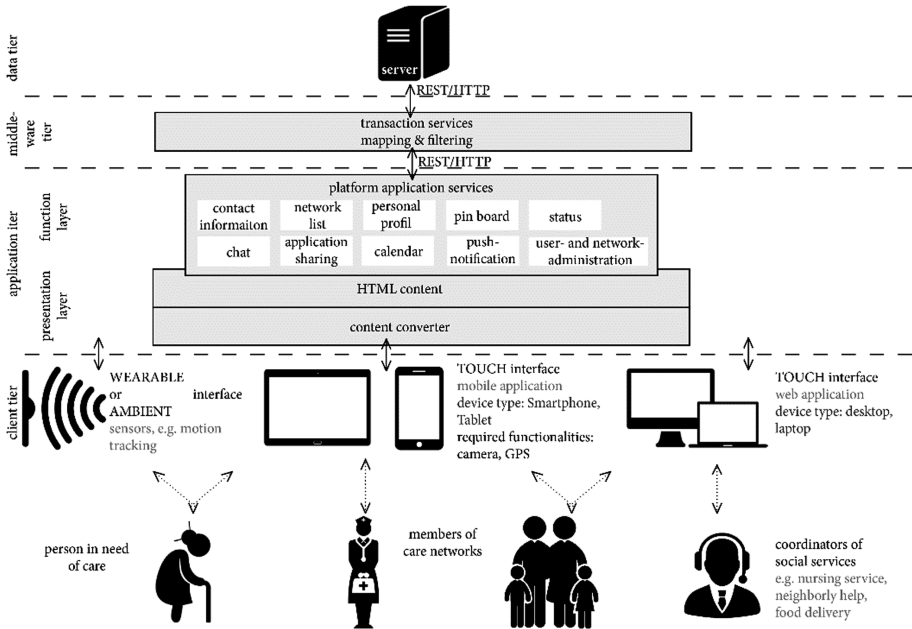


Fig. 1. Presentation of the collected requirements in a four-tier architecture system design

Another often discussed topic is the integration of sensors, such as life-sign buttons, medication dispensers or motion detectors in the apartment of the person in need of care, which provide additional information and convey a comforting feeling for the whole network and especially the care responsible.

Whereas [11] set to a complete proprietary development, in the design process the need to support different communication channels and optimally integrate these to one stream, to reach out to all actors, was identified as a further requirement for the application. Push-notifications generate instant awareness and reduce the impact of media breaks. Against this backdrop, the solution should be set up as a mixture of proprietary development (e.g. chat, pin board) and external app launcher (e-mail, phone, internet browser).

A frequently mentioned aspect is the internet connection. In that regard, rural more than urban participants complain about poor internet connections. Especially the case studies showed that a mobile support tool must include offline functionalities. A separated system architecture – containing a middleware tier for transaction services, mapping and filtering, as well as a data tier – ensures low data traffic and fast loading times. A synchronization of all essential information in the application tier guarantees offline services (c.f. Fig. 1).

The mock-up presented in Fig. 2 has already been transformed into a first prototype and is ready for a first field test. The application tier is implemented using the frontend frameworks IONIC and Angular. Due to limited manpower, this first prototype does not communicate with a middleware but directly with the server.

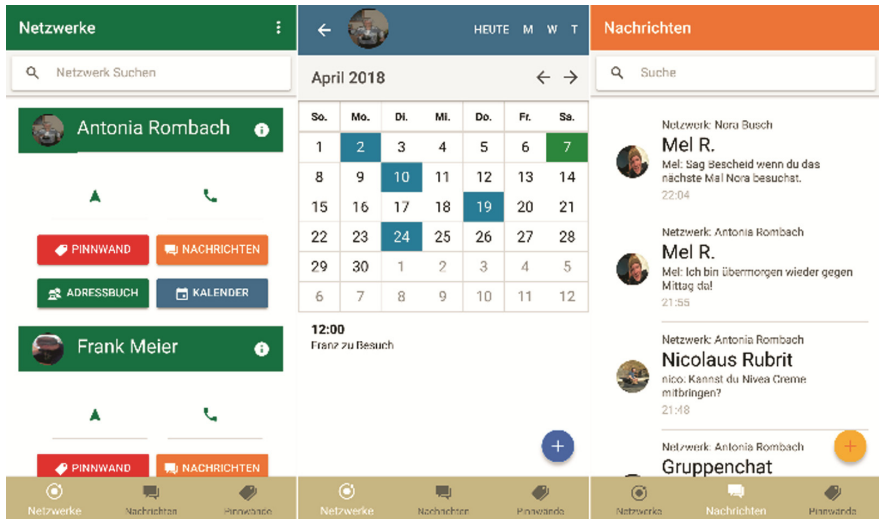


Fig. 2. Mock-ups of the concept created with IONIC. (left) Start view which presents an overview of all networks (middle) network specific calendar (right) message overview

6 Discussion and Conclusion

The presented concept is particularly geared towards supporting communication and cooperation in care networks. The designed system does not merely constitute a standard communication tool, but rather offers participants a comprehensive overview over the different roles within the network, enables to share case-related information via a pin board, facilitates scheduling and allows transferring case-related appointments to private calendars. Hence, apart from the goal to link the individual actors within networks, it builds awareness for a caring community.

Care Recipients' Perspectives: Care recipients of mental health should be granted access to the collaboration platform to enable them to participate. As described in [14], with this tool it is likewise possible that helpers request for help themselves if necessary. The authors included potential care recipients (senior neighborhood helpers) and family care givers in the conception of the platform. Mentally confused care recipients are according to the authors opinion not considered target users and were hence not included in the needs assessment.

Professional Caregivers' Perspectives: The proposed technical support requires a further organizational development and a changed understanding of roles of the professional caregivers. The experiences of the workshops showed that there is no theoretical background or basic understanding of the topic. In order to enhance this essential understanding of the different goals and tasks of primary systems and collaboration tools, a high level of explanation is needed. Different project activities showed that improving this understanding and building awareness is quite time-consuming and needs

persuasion. It is important to comprehensively explain that collaboration tools do not double the work, but rather render time demanding phone calls unnecessary, offer a means for informal interchange and support information sharing, which is especially useful in case of an emergency. Nevertheless, in view of a seamless electronic communication, software interfaces to standard primary documentation software tools should be considered in the implementation of the presented concept.

Data Security, Roles and Permissions: The concept is designed to enhance collaboration in care-mix settings. The documentation of health-sensitive data which may stand in conflict with data security laws is not intended. Nevertheless, it cannot be ruled out that future users may misuse the tool. So far two roles (network administrators and members) with different access restrictions are implemented. More complex role and access models were evaluated as restrictive and confusing.

Design – One for All or All One’s Own? While it is generally agreed among the stakeholder groups that the collaborative functionalities of a support platform are necessary, concept workshops showed quite diverse profiles of requirements per group. Other study activities, however, showed that it is not necessary to have one single design for all groups neither. According to the results, three designs are appropriate to cover the needs of all stakeholders. One user interface should be designed for single network users (less technically-versed persons). A second application should be designed for members of multiple care-mix networks. (e.g., nurses, family doctors, neighborhood helpers) (c.f. Fig. 1). The third design should be set up for professional actors like managers of care services or organized neighborhood help. These stakeholders usually hold special roles as advisors and contact persons, which requires a different overview of the networks, persons, capabilities and resources.

Regardless of this, in order to find acceptance, the design must be appealing and meet the high standards of universal collaboration software.

Further Limitations of the Presented Work and Future Activities: In total, 35 participants were surveyed during the study. The results therefore are only of qualitative nature. However, the inclusion of findings from current literature significantly increases the relevance of the presented results. At this point it must be noted that it is very difficult to reach the already stressed target group at all.

Future activities in this project include field tests of the prototype in several care-mix settings. Each care-mix team will be closely supported by workshops and surveys and accompanied for at least six months. Thereby, a special focus will be on analyzing the influencing factors for a successful implementation.

References

1. Triantafillou, J.: Supporting Family Carers of Older People in Europe-The Pan-European Background Report (2005)
2. Bäuerle, D., Scherzer, U.: Zukunft Quartier – Lebensräume zum Älterwerden. Themenheft 1: Hilfe-Mix – Ältere Menschen in Balance zwischen Selbsthilfe und (professioneller) Unterstützung (2009). http://www.netzwerk-song.de/fileadmin/user_upload/Themenheft1.pdf

3. Görres, S., Seibert, K., Stiefler, S.: Perspektiven zum pflegerischen Versorgungsmix. Pflege-report 2016, pp. 3–14. Schattauer GmbH, Stuttgart (2016)
4. Pinelle, D., Gutwin, C.: A groupware design framework for loosely coupled workgroups. In: Gellersen, H., Schmidt, K., Beaudouin-Lafon, M., Mackay, W. (eds.) ECSCW 2005, pp. 65–82. Springer, Dordrecht (2005). https://doi.org/10.1007/1-4020-4023-7_4
5. Szabo, E.: Organisationskultur-Forschung im Überblick. In: Organisationskultur und Ethnographie, pp. 1–46. Deutscher Universitätsverlag, Wiesbaden (1998)
6. Muller, M.J., Kuhn, S.: Participatory design. *Commun. ACM* **36**, 24–28 (1993)
7. Mayring, P.: Qualitative inhaltsanalyse. In: *Handbuch Qualitative Forschung in der Psychologie*, pp. 601–613. VS Verlag für Sozialwissenschaften, Wiesbaden (2010)
8. Hutchinson, H., Mackay, W., Westerlund, B., Bederson, B.B., Druin, A., Plaisant, C., Beaudouin-Lafon, M., Conversy, S., Evans, H., Hansen, H.: Technology probes: inspiring design for and with families. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 17–24 (2003)
9. Rosson, M.B., Carroll, J.M.: Scenario-based design. In: Sears, A., Jacko, J.A. (eds.) *Human-Computer Interaction*, pp. 145–162. CRC Press, Boca Raton (2009)
10. Fitzpatrick, G., Ellingsen, G.: A review of 25 years of CSCW research in healthcare: contributions, challenges and future agendas. *Comput. Support. Coop. Work* **22**, 609–665 (2012)
11. Bossen, C., Christensen, L.R., Grönvall, E., Vestergaard, L.S.: CareCoor: augmenting the coordination of cooperative home care work. *Int. J. Med. Inform.* **82**, e189–e199 (2013)
12. Moser, C., Krischkowsky, A., Neureiter, K., Tscheligi, M.: Mediating informal care online: findings from an extensive requirements analysis. *Interact. Des. Archit. J.* **24**, 33–48 (2015)
13. Abou Amsha, K., Lewkowicz, M.: Shifting patterns in home care work: supporting collaboration among self-employed care actors. In: De Angeli, A., Bannon, L., Marti, P., Bordin, S. (eds.) *COOP 2016: Proceedings of the 12th International Conference on the Design of Cooperative Systems*, 23–27 May 2016, Trento, Italy, pp. 139–154. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-33464-6_9
14. Schorch, M., Wan, L., Randall, D.W., Wulf, V.: Designing for those who are overlooked - insider perspectives on care practices and cooperative work of elderly informal caregivers. In: *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work and Social Computing, CSCW 2016*, pp. 785–797 (2016)
15. Bødker, S., Grönvall, E.: Calendars: time coordination and overview in families and beyond. In: Bertelsen, O., Ciolfi, L., Grasso, M., Papadopoulos, G. (eds.) *ECSCW 2013: Proceedings of the 13th European Conference on Computer Supported Cooperative Work*, 21–25 September 2013, Paphos, Cyprus, pp. 63–81. Springer, London (2013). https://doi.org/10.1007/978-1-4471-5346-7_4
16. Ganoë, C.H., Somervell, J.P., Neale, D.C., Isenhour, P.L., Carroll, J.M., Rosson, M.B., McCrickard, D.S.: Classroom BRIDGE: using collaborative public and desktop timelines to support activity awareness. In: *Proceedings of the 16th Annual ACM Symposium on User Interface Software and Technology*, pp. 21–30. ACM Press, New York (2003)
17. Bastide, R., Bardy, P., Borrel, B., Boszodi, C., Bouet, M., Gani, K., Gayraud, E., Gourc, D., Lamine, E., Manenq, P.H., Schneider, M., Toumani, F.: Plas ‘O’ Soins: a software platform for modeling, planning and monitoring homecare activities. *IRBM* **35**, 82–87 (2014)



Identifying Suitable Representation Techniques for the Prioritization of Requirements and Their Interdependencies for Multiple Software Product Lines

Stephanie Lewellen^(✉) and Markus Helfert

Dublin City University, Glasnevin, Dublin 9, Ireland
stephanie.lewellen2@mail.dcu.ie, markus.helfert@dcu.ie

Abstract. Software requirements typically do not exist independently of each other, rather most requirements have some type of dependency on another requirement [4]. For companies developing software products, which depend on each other, in so-called multiple software product lines (SPLs), systematic requirements management, including consideration for prioritization and interdependencies, is a time-consuming and convoluted task. Representation techniques for complex requirements can convey critical requirements interdependency information to make prioritization of requirements quicker and more accurate [1]. Based on reviewing the foremost literature, this paper identifies the representation techniques for requirements management which are most suitable for multiple software product lines (SPLs).

Keywords: Requirements · Prioritization · Multiple software product lines
Interdependencies

1 Introduction

Software release planning is a critical decision-making process which aims to find an optimal subset of software requirements, in which the stakeholders are satisfied while the resource and timeline constraints are met [2, 3]. Software release planning is one of the most important and complex tasks within the practices of requirements engineering, because requirements usually have many dependencies and it is not possible to select requirements based on individual priority alone [4].

One of the catalysts to increased complexity in requirements dependencies is modern software product line engineering, which capitalizes on the recycling of software between software products [5]. In the simplest form, the software product line (SPL) has one common code base, but with two or more higher level code additions resulting in their own software products [5, 6]. In recent years, with even more variation, the concept of SPLs has been extended to offer more software product variation in the form of multiple SPLs [5, 6].

For multiple SPLs, requirements interdependencies which influence the cost of development and the value to the stakeholder could be of particular interest, because if

they are not met, the original benefit of diversifying the software product portfolio has not been completely realized.

This paper considers the representation techniques in the prioritization of requirements which have interdependencies between SPLs. Focus is given to the interdependencies of requirements which are most critical to market-driven software [1] – requirements related to value and cost. It starts with an overview of the current research in the areas of multiple SPLs and requirements management, including prioritization of requirements and requirements interdependencies. Following the overview is a summary of five representation techniques for the prioritization of requirements and their interdependencies. Finally, the paper concludes with the authors' assessment of the five representation techniques based on requirement prioritization, requirement interdependencies, and the consideration of cost/value interdependencies between SPLs.

2 Overview

This overview delves deeper into the subject areas and foremost research of multiple SPLs and managing software requirements, including the prioritization of requirements and assessment of requirement interdependencies.

2.1 Multiple Software Product Lines (SPLs)

Much consideration has been given to requirements management for single software products, but far less consideration has been given to requirements between software product lines [7] or to particular domains [8].

A software product line (SPL) is a set of software products sharing a set of common features but containing variation points [9]. One of the advantages to SPLs is the reduced cost of development and testing, with an increased opportunity to address different stakeholder groups [7]. An example for a single SPL, given by Rosenmüller et al. [10], is mail client software which relies on a common mail framework. Variations of the mail client software, to support different protocols, for example, would all rely on the same common mail framework [10].

Multiple SPLs, in comparison, are composed of many interconnected subsystem versions and variants [6]. Multiple SPLs commonly refers to vertically tiered software stacks, with application SPLs and infrastructure SPLs [5], but can also refer to distributed SPLs, as in the case of sensor software [10]. An example of vertically tiered multiple SPLs would be individual application software product lines that all rely on the same database platform, which is itself a software product line [5, 6] (Fig. 1).

The analysis that follows in this paper focuses on requirements representation techniques and their application to requirements prioritization and requirements interdependencies between multiple SPLs, specifically vertically tiered SPLs where the upper tier (application) software requires software functionality from the lower (infrastructure) tiers.

Furthermore, the focus in this paper is on mature products on the market, with development approach that is incremental and market-driven, as in the industrial survey conducted by Carlshamre et al. [1]. The focus has been restricted to market-driven

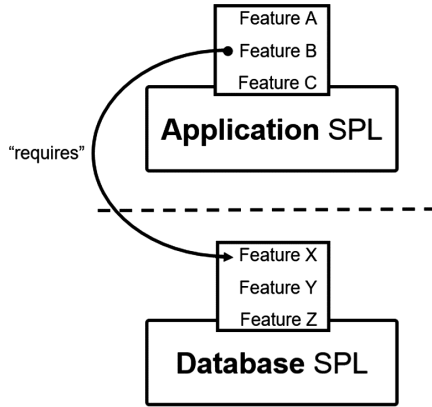


Fig. 1. Dependency model between an application SPL and an infrastructure SPL [5]

software, because most requirements interdependencies for this type of software are of the “cost/value” type [1].

2.2 Managing Software Requirements: Prioritization and Interdependencies

It is not possible to address all functional and non-functional requirements in the next software release, due to constraints from budget, resources or time [7]. Therefore, a prioritization of requirements should be applied [8].

Prioritization Consideration. There are a number of requirements prioritization techniques [11, 12] and the research to the application of these processes to the software release planning process is extensive. However, even the most suitable processes for complex requirement prioritization, like the analytic hierarchy process (AHP), do not consider the interdependency of requirements [8, 12].

Interdependency Consideration. According to Dahlstedt et al., “*despite the need for and potential benefits of systematically taking requirements interdependencies into account, there is little research invested in this topic and more is needed*” [13]. The existing processes [12] provide potentially inaccurate weighting of requirements due to these shortcomings [14].

Carlshamre et al. [1] noted in the result of their 2001 industrial survey that approximately 20% of the requirements assessed in their survey had 75% of the total requirement dependencies. In order to reduce an otherwise time-intensive process of assessing all requirements for interdependencies, it would make sense to identify the requirements with obvious interdependencies first, and represent them visually so that interdependency information could be inferred quickly [1].

Requirement Interdependency Types and Their Suitability to Multiple SPLs.

Before we can discuss the most suitable representation for requirements between multiple SPLs, it is important to (1) describe the types of requirements interdependencies

which have been identified in the foremost literature, (2) identify the requirements interdependency types which are most critical to the scenario that one SPL has a requirement dependency to another SPL.

Tables 1 and 2 provide the most referenced sources for requirement interdependency categories [1, 13], and are in general alignment with each other on interdependency types.

Table 1. Interdependency categorization by Dahlstedt et al. [13]

Interdependency categorization	Interdependency type
Structural	Require
	Explain
	Similar to
	Conflict with
	Influences
Cost/Value	Increase/Decrease cost
	Increase/Decrease value

Table 2. Interdependency types by Carlshamre et al. [1]

Priority, lowest number takes precedence	Interdependency type, where R = requirement	Meaning, where R = requirement
1	R ₁ AND R ₂	R ₁ requires R ₂ to function, and R ₂ requires R ₁ to function.
2	R ₁ REQUIRES R ₂	R ₁ requires R ₂ to function, but not vice versa.
3	R ₁ TEMPORAL R ₂	Either R ₁ has to be implemented before R ₂ or vice versa.
4	R ₁ CVALUE R ₂	R ₁ affects the value of R ₂ for a customer. Value can be either positive or negative.
4	R ₁ ICOST R ₂	R ₁ affects the cost of implementing R ₂ . Value can be either positive or negative.
5	R ₁ OR R ₂	Only one of {R ₁ , R ₂ } needs to be implemented.

These requirement interdependency types provide the basis for a more complex analysis of how interdependencies, sometimes also called requirements relations, are considered in representation techniques for multiple SPLs. In Fig. 2, the difference between a requirement interdependency within one SPL and an interdependency between multiple SPLs is depicted. While most of the foremost literature and representation techniques consider the simple case of a requirement requiring another requirement within the same SPL, the case we want to consider is how the foremost representation techniques are suitable to the more complex case that there is a requirement interdependency between product lines.




Fig. 2. Requirement interdependency within one SPL as opposed to between multiple SPLs

Because each SPL is, in its software productization, independent of the other one, the separate product lines have the potential to have different release timelines, market segments, share of company operational revenue, etc. [5]. Therefore, we can identify some existing requirement interdependency types to be more critical for requirements between SPLs.

We identify the “requires” type to be critical due to the definition of vertically tiered multiple SPLs. If a requirement between SPLs exist, it will always be a requirement from the top-most-level SPL of the SPL tier beneath it. Additionally, any requirement between SPLs most likely also has a “cost/value” interdependency type due to the nature of SPLs, which are designed to increase the value of software options at reduced overall development cost [1].

With this research, we aim to answer: which of the foremost techniques for representing the prioritization and interdependencies of requirements within one SPL is the most suitable for representing the more complex requirements interdependencies between multiple SPLs?

3 Representation Techniques for Requirements Prioritization and Requirements Interdependencies

We have identified the following representation techniques as part of a systematic literature review to be the most commonly referenced.

3.1 Directed Graph (Digraph) Representation

A directed graph (digraph) represents requirements as shapes connected by arrows. By differentiating requirements prioritization and interdependencies by color, line type, and size of the shapes and arrows, the digraph is able to represent complex relations between the requirements.

A practical example of a digraph requirements representation was made by Carlshamre et al. [1] when they took an industrial survey where requirements managers from five software organizations were asked to perform pairwise assessment on requirements, only considering priority. A pairwise assessment was performed using a spreadsheet designed by Carlshamre et al. and ensured that all requirements were compared with each of the other requirements.

The requirements managers were then asked to identify interdependencies between the requirements. They were also asked to give each interdependency a simple certainty rating (possibly-probably-positively). The interdependency types available in the identification are shown in Table 2 and include the type REQUIRES (dependency), CVALUE (customer value), and ICOST (increases cost).

In addition to identifying interdependencies with certainty ratings, Carlshamre et al. also define a hierarchy between interdependency types for the case that more than one relationship is identified between two requirements. The priority (or hierarchy) of the relationships is also shown in Table 2 with the REQUIRES interdependency type assigned priority 2, and both CVALUE and ICOST assigned priority 4. In the case that CVALUE and ICOST have a conflict, they have to be traded off against each other [1].

Using this data, Carlshamre et al. created digraphs of the requirements priorities and the requirements interdependencies for each software organization. By representing the requirements, their priorities, and their interdependencies by objects and arrows, it is possible to draw important conclusions just from a glance at the digraph. However, the authors of this paper observe that the graphical representation itself does not convey some of the more sophisticated data collected, like the certainty rating for the interdependencies.

3.2 Metamodel Ontology Representation

A metamodel ontology is an appropriate method for describing and visualizing requirements, their prioritization and their interdependencies because the model supports distilling the relationships between requirements to their base elements [7, 9, 15, 16].

Due to the flexibility of the metadata in a metamodel ontology, it is possible to rate the priority on a scale. The common structural interdependencies can be modeled, which include a “requires” interdependency type, however in the foremost literature, there is no example of modelling “cost/value” interdependencies in a metamodel ontology.

3.3 Software Requirements Catalog (SRC) Method

The software requirements catalog (SRC) is a method for collecting and considering software requirements for reuse instead of considering individual software features or

components for reuse. The creation of the SRC includes a classification phase, where the functionality of the requirement to fulfill the goal – the reason for the existence of the project – is described. The description includes a prioritization rating which reflects the suitability of the requirement to the project goal [17–19].

There is a qualitative high-medium-low rating scale for evaluating the priority of the requirement to fulfill the project goal. However, the definition of a “priority” in the SRC context is a variation on the definition of priority we have discussed previously. In the context of SRC, priority is a rating of the requirement in only the parameter of its suitability to fulfill the overall project purpose, and not of its overall criticality [17].

The requirement constraints and dependencies with other software projects are defined and refined in order to continuously update the requirements catalog [17]. The tracking of requirements interdependencies, also called traceability, seems to consider simple relationships, such as the “requires” structural interdependency between projects with common requirements. However, there is no mention of more complex requirements interdependencies between projects like “cost/value” interdependencies.

3.4 Fuzzy Graph Representation

Requirements interdependencies are considered fuzzy relations because the strengths of the dependencies can vary greatly [20]. Mougouei et al. model the influence of value-related requirements interdependencies using fuzzy graphs, which consider the uncertainty of the dependency relations.

Although there are also fuzzy representation techniques available for the prioritization of requirements, including the fuzzy AHP technique [11], none of those techniques take the interdependency of requirements into consideration [12, 13].

3.5 Cost-Value Diagram Representation

The cost-value approach to requirements management involves a pair-wise comparison on requirements in two dimensions: the requirement value and requirement cost [21]. Based on the results, the requirements are depicted on a graph with an axis for value and axis for cost and two delineations to fence off the high-value/low-cost requirements, mid-value/mid-cost requirements, and low-value/high-cost requirements. Karlsson et al. [21] developed a support tool to plot the scored requirements and were also able to take simple structural interdependencies into account, including the “requires” interdependency type [21].

4 Criteria and Comparison of Requirement Prioritization and Interdependency Representation Techniques Suitable for Multiple SPLs

In the following section, we present the criteria we used to assess the requirement prioritization and interdependency representation techniques suitable for multiple SPLs. We then compare the overall rating for the representation types discussed in Sect. 3.

4.1 Criteria for Rating

A simple (SMART) scoring technique has been applied to each of the dimensions for analysis [22].

Priority Consideration. Each representation type in Sect. 3. was evaluated against the sophistication of requirement priority consideration on a decimal scale from zero to one using the criteria in Table 3. A rating of zero corresponds to an absence of consideration. A rating of 0.5 corresponds to a simple scale priority rating, where requirements are given a standalone rating. A rating of 1.0 corresponds to a comparative prioritization where the requirements are compared to one another and then receive a relative priority rating.

Table 3. Criteria for priority consideration rating [22]

Rating	Criteria
0	Absence of priority consideration
0.5	Simple scale priority rating (ex. low, medium, high)
1.0	Relative prioritization using requirement comparison

Interdependency Consideration. Each representation type in Sect. 3 was evaluated against the sophistication of requirement interdependency consideration on a decimal scale from zero to one using the criteria in Table 4. A rating of zero corresponds to an absence of consideration. A rating of 0.5 corresponds to a simple interdependency tracking, where it is represented that requirements are linked. A rating of 0.75 corresponds to an interdependency consideration with a certainty rating. A rating of 1.0 corresponds to a complex interdependency consideration, where, for example, multiple types of interdependencies are represented.

Table 4. Criteria for interdependency consideration rating [22]

Rating	Criteria
0	Absence of interdependency consideration
0.5	Simple interdependency tracking (traceability)
0.75	Interdependency consideration with certainty rating
1.0	Complex interdependency consideration

Suitability for Multiple SPLs. Each representation type in Sect. 3 was evaluated against the suitability of the type for multiple SPLs on a decimal scale from zero to one using the criteria in Table 5. Specifically, the suitability criteria refer to the interdependency types “requires” and “cost/value”, which play a critical role in the assessment of requirements interdependencies between SPLs. A rating of zero corresponds to an absence of consideration for even the most basic structural interdependency type, “requires”. A rating of 0.5 corresponds to a simple consideration for either “requires” or “cost/value” interdependency types. A rating of 0.75 corresponds to either consideration for both “requires” and “cost/value” interdependency types or an in-depth consideration of either one. A rating of 1.0 corresponds to a representation type that allows for

requirement consideration for requirements from external SPLs, which bring potentially their own “cost/value” and “requires” requirements.

Table 5. Criteria for multiple SPL suitability rating [22]

Rating	Criteria
0	Absence of consideration for interdependency type “requires” and “cost/value”
0.5	Simple consideration for interdependency type “requires” or “cost/value”
0.75	Consideration for “requires” and “cost/value”, or in-depth handling for one or the other
1.0	Additional parameters for “cost/value” to consider one or more dependent SPLs outside of the assessed SPL itself

4.2 Comparison of Representation Techniques

The following is a comparison of the previously discussed requirements prioritization and requirements interdependencies techniques using the criteria specified in the previous subsection.

The summary of the findings is that even though some of the techniques take into account all three criteria areas, prioritization, interdependencies, and multiple SPL requirements, none of them offer a complete solution to requirements prioritization with interdependencies between multiple SPLs.

The cumulative total ratings from Table 6 are visualized as a multi-dimensional graph in Fig. 3 to add a qualitative perspective to the quantitative ratings.

Table 6. Rated comparison of representation techniques

Technique	Priority consideration	Interdependency consideration	Suitability for multiple SPLs	Cumulative total rating
Digraph	1.0	0.75	0.75	2.5
Metamodel ontology	0.5	0.5	0.5	1.5
SRC	0.5	0.5	0.5	1.5
Fuzzy graph	0	1.0	0.75	1.75
Cost-value diagram	1.0	0.75	0.75	2.5

It was determined based on the given criteria that the SRC method and metamodel ontology both had a simple prioritization method and a simple interdependency tracking method. They both also had no mention of the “cost/value” requirements interdependency type, which is critical to a thorough representation of requirements between SPLs. The metamodel ontology representation rates slightly higher on the graph for multiple SPL suitability because it has greater potential for further extensibility through metadata. The SRC method limits flexibility to requirements interdependency handling between SPLs because the requirement scoring is from the perspective of the requirement to fulfill a specific project, and is therefore better suited to smaller software projects with fewer diverse stakeholders.

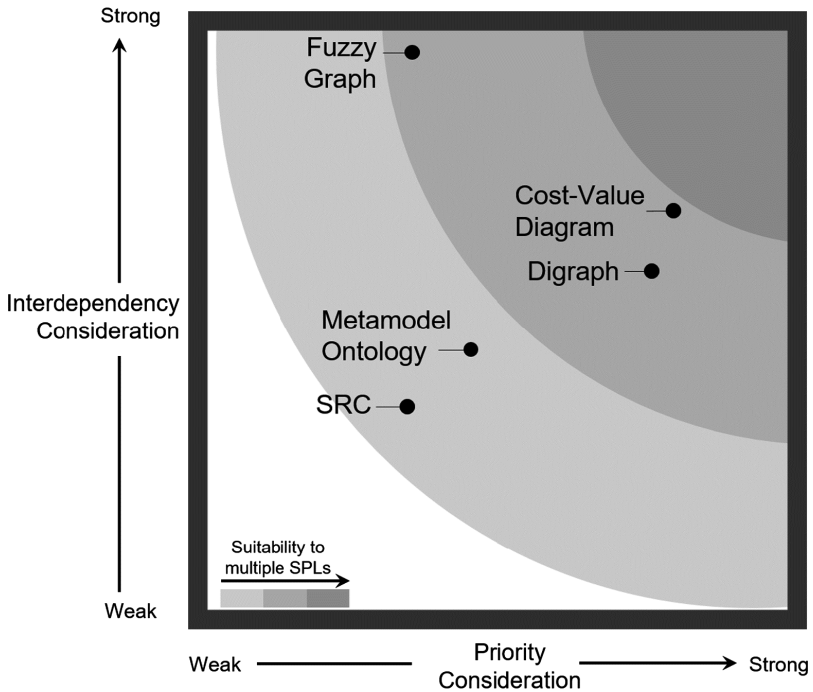


Fig. 3. Multi-dimensional graph of rated representation techniques

Slightly more suited for requirements between SPLs is the fuzzy graph representation, which gives very detailed information about requirements interdependencies and their uncertainties, but does not take prioritization into consideration at all.

Most suitable to requirements between SPLs are the directed graph (digraph) model and the cost-value diagram. Both are capable of distilling the required requirements interdependency types in a compact and straightforward way. Both could potentially be extended to represent relationships between SPLs. However, the cost-value diagram adds an additional perspective of potential investment areas (ex. high-value/low-cost requirements), which could be valuable if extended to multiple SPL requirements support.

5 Conclusion

The assessment of the representation techniques in Sect. 4.2 Comparison of Representation Techniques shows that the cost-value diagram and the digraph representations are the closest to being suitable for multiple SPLs. More research is necessary into these representation techniques, specifically in how they could be extended to more accurately represent requirement interdependencies and the priorities thereof for multiple SPLs.

References

1. Carlshamre, P., Sandahl, K., Lindvall, M., Regnell, B., och Dag, J.: An industrial survey of requirements interdependencies in software product release planning. In: Proceedings of the Fifth IEEE International Symposium, Requirements Engineering (2001)
2. Mougouei, D., Powers, D.M., Moeini, A.: Dependency-aware software release planning. In: IEEE/ACM 39th International Conference Software Engineering Companion (ICSE-C), 2017 (2017)
3. Bagnall, A.J., Rayward-Smith, V.J., Whittle, I.M.: The next release problem. *Inf. Softw. Technol.* **43**(14), 883–890 (2001)
4. Carlshamre, P., Regnell, B.: Requirements lifecycle management and release planning in market-driven requirements engineering processes. In: 11th International Workshop Database and Expert Systems Applications, Proceedings (2000)
5. Schirmeier, H., Spinczyk, O.: Challenges in software product line composition. In: 42nd Hawaii International Conference, System Sciences, HICSS 2009 (2009)
6. Damiani, F., Schafer, I., Winkelmann, T.: Delta-oriented multi software product lines. In: Proceedings of the 18th International Software Product Line Conference, vol. 1 (2014)
7. Soomro, S., Hafeez, A., Shaikh, A., Musavi, S.H.A.: Ontology based requirement interdependency representation and visualization. In: Shaikh, F.K., Chowdhry, B.S., Zeadally, S., Hussain, D.M.A., Memon, A.A., Uqaili, M.A. (eds.) *IMTIC 2013*. CCIS, vol. 414, pp. 259–270. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-10987-9_24
8. Berander, P., Andrews, A.: Requirements prioritization. In: Aurum, A., Wohlin, C. (eds.) *Engineering and Managing Software Requirements*, pp. 69–91. Springer, Heidelberg (2005). https://doi.org/10.1007/3-540-28244-0_4
9. Sellier, D., Mannion, M., Mansell, J.X.: Managing requirements inter-dependency for software product line derivation. *Requirements Eng.* **13**(4), 299–313 (2008)
10. Rosenmüller, M., Siegmund, N., Kästner, C., ur Rahman, S.: Modeling dependent software product lines. In: Proceedings of the GPCE Workshop on Modularization, Composition and Generative Techniques for Product Line Engineering (McGPLE) (2008)
11. Achimugu, P., Selamat, A., Ibrahim, R., Mahrin, M.N.: A systematic literature review of software requirements prioritization. *Inf. Softw. Technol.* **56**(6), 568–585 (2014)
12. Kahn, J., Rehman, I.: Comparison of requirement prioritization technique to find the best prioritization technique. *Modern Educ. Comput. Sci.* **7**(11), 53–59 (2015)
13. Dahlstedt, A., Persson, A.: Requirements Interdependencies: State of the Art and Future Challenges. In: Aurum, A., Wohlin, C. (eds.) *Engineering and Managing Software Requirements*, pp. 95–116. Springer, Heidelberg (2005). https://doi.org/10.1007/3-540-28244-0_5
14. Zhang, H., Li, J., Zhu, L., Jeffery, R., Liu, Y., Wang, Q., Li, M.: Investigating dependencies in software requirements for change propagation analysis. *Inf. Softw. Technol.* **56**(1), 40–53 (2014)
15. Goknil, A., Kurtev, I., van den Berg, K., Veldhuis, J.: Semantics of trace relations in requirements models for consistency checking and inferencing. *Softw. Syst. Model.* **10**(1), 31–54 (2011)
16. Saeki, M., Kaiya, H.: On Relationships among models, meta models, and ontologies. In: Proceedings of the Proceedings of the 6th OOPSLA Workshop on Domain-Specific Modeling (DSM 2006) (2006)
17. Pacheco, C., Garcia, I., Calvo-Manzano, J., Arcilla, M.: Reusing functional software requirements in small-sized software enterprises: a model oriented to the catalog of requirements. *Requirements Eng.* **22**(2), 275–287 (2017)

18. Robertson, S., Robertson, J.: *Mastering the Requirements Process*, 6th edn. Addison-Wesley, Munich (2013)
19. Wiegers, K.: First things first: prioritizing requirements. *Software Dev.* **7**(9), 48–53 (1999)
20. Mougouei, D., Powers, D.: Modeling and selection of interdependent software requirements using fuzzy graphs. *Int. J. Fuzzy Syst.* **19**(6), 1812–1828 (2017)
21. Karlsson, J., Olsson, S., Ryan, K.: Improved practical support for large-scale requirements prioritising. *Requirements Eng.* **2**(1), 51–60 (1997)
22. Valiris, G., Chytas, P., Glykas, M.: Making decisions using the balanced scorecard and the simple multi-attribute rating technique. *Perform. Meas. Metrics* **6**(3), 159–171 (2005)

Author Index

- Abdelsalam, Hisham M. 206
Afifi, Nehal 206
Awad, Ahmed 206
- Batoulis, Kimon 236
Behrens, Dennis 58, 265
Böhm, Markus 295
Borbinha, José 102
Buijs, J. C. A. M. 163
Burnaev, Evgeny 347
- Cimmino, Andrea 307
Corchuelo, Rafael 307
- Decker, Stefan 384
Derouet, Maximilian 30
Dixit, P. M. 163
- Fani Sani, Mohammadreza 115
- Gevorkjan, Gor Davidovic 148
Ghabri, Rachaa 3
- Habrich, Thilo 359
Hein, Andreas 295
Helfert, Markus 412
Hellingrath, Bernd 359
Hewelt, Marcin 223
Hirmer, Pascal 3
Holzleitner, Manuel 278
- Ignatiev, Vladimir 347
- Kaufhold, Anna 58
Khairova, Nina 333
Kirilov, Lyubomir 321
Klein, Achim 321
Klohs, Katharina 45
Knackstedt, Ralf 58, 265
Konstantinidis, Ioannis 384
Kremer, Helmut 295
Krogstie, John 135
- Kuehnel, Stephan 176
Kunze, Christophe 400
Kuralai, Mukhsina 333
- Leukel, Joerg 321
Lewellen, Stephanie 412
Lewoniewski, Włodzimierz 333
Leyh, Christian 86
Lindner, Dominic 86
Linn, Christian 30
Lukasiewicz, Wojciech 15
- Maggi, Fabrizio M. 190, 371
Mikalef, Patrick 135
Milani, Fredrik 371
Mitschang, Bernhard 3
- Nalbach, Oliver 30
Nikaj, Adriatik 223
Novikov, German 347
- Orken, Mamyrbayev 333
- Paschke, Adrian 15
Peristeras, Vassilios 384
Potapov, Georgy 347
Proença, Diogo 102
- Reijers, Hajo A. 190
Renyi, Madeleine 400
Riekert, Martin 321
Ritter, Daniel 278
Rossmann, Alexander 71
- Sackmann, Stefan 176
Sandkuhl, Kurt 45
Schoormann, Thorsten 58, 265
Schunselaar, Dennis M. M. 190
Seyffarth, Tobias 176
Siaminos, Georgios 384
Slaats, Tijs 190

Teuteberg, Frank 148, 253, 400

Timm, Felix 45

Timplalexis, Christos 384

Todor, Alexandru 15

Tönnissen, Stefan 253

Trekin, Alexey 347

van de Wetering, Rogier 135

van der Aalst, Wil M. P. 115, 163, 190

van Zelst, Sebastiaan J. 115

Vejseli, Sulejman 71

Verbeek, H. M. W. 163

Wagner, Carolin 359

Warnecke, Danielle 148

Węcel, Krzysztof 333

Werth, Dirk 30

Weske, Mathias 223, 236

Zervas, Panagiotis 384