

A Taxonomic Framework for Social Machines

Paul Smart, Elena Simperl, and Nigel Shadbolt

1 Introduction

Within the context of the World Wide Web, we have witnessed the emergence of a rich range of technologies that support both collaboration and distributed processing. Applications such as Wikipedia, for instance, have demonstrated the power and potential of the Web to facilitate the pooling of geographically dispersed knowledge assets. The result has been the creation of the world's largest online encyclopedia, available for free in more than 200 languages for everyone to access and use. Similar success stories can be found in various other domains. Projects such as Galaxy Zoo,¹ for example, have shown how collective intelligence can be used to inform scientific inquiry, while initiatives such as Ushahidi² have played a crucial role in emergency management situations worldwide. These three Web-based systems are representative of a general trend which is characterized by the use of Web-based technologies to enable a wide range of activities that rely on a combination of decentralized human activity and computational processing. This trend has been reflected in research efforts across a variety of areas, including social computing [32], human computation [34], crowdsourcing [11], and collective intelligence [5]. It has also given rise to a variety of new concepts, such as the 'social computer' [38], the 'social operating system' [35] and 'social machines' [18]. This chapter focuses on the last term in this list: the concept of social machines. The term 'social machine' was first used in a Web context by Berners-Lee and

¹<http://www.galaxyzoo.org/>.

²<http://www.ushahidi.com/>.

P. Smart (✉) • E. Simperl • N. Shadbolt
Electronics & Computer Science, University of Southampton, Southampton, SO17 1BJ, UK
e-mail: ps02v@ecs.soton.ac.uk; e.simperl@soton.ac.uk; nrs@ecs.soton.ac.uk

Fischetti [3], and it has since grown in popularity to the point where it is now the focus of large-scale research programs, such as the EPSRC's SOCIAM initiative,³ the subject of a multitude of academic papers (e.g., [18, 26, 27, 30, 41, 43, 50]) and the basis for a workshop series at the World Wide Web conference.

In spite of the growing interest in social machines, however, there is little consensus, at the present time, as to what the term 'social machine' actually means. In addition, the scientific community seems to have only a very narrow understanding as to what kinds of social machines actually exist. In order to make progress in these areas, we attempt to provide a working definition of the social machine concept that builds on the ideas put forward by Berners-Lee and Fischetti [3]. We also introduce a taxonomic framework for social machines that features a set of dimensions along which all social machines are deemed to vary. This work extends the results of an earlier study, reported by Shadbolt et al. [43], which used knowledge elicitation techniques to generate an initial set of dimensions. The work reported in the current chapter differs from this earlier body of work in two ways. Firstly, the dimensions from the earlier study have been refined and enriched following discussions with members of both the computer science and social science communities. Secondly, the current framework features a complete set of characteristics for each dimension. These characteristics specify the 'values' that each social machine takes with respect to each of the dimensions in the framework (see Sect. 4).

Together, the effort to provide a definition for social machines and the effort to develop a taxonomic framework mark an important step in terms of our attempt to understand the emerging, interdisciplinary research field of social machines. The effort to provide a definition of social machines is crucial because in the absence of an ability to say what social machines are it becomes difficult to know where to focus research and engineering efforts. The lack of a definition also complicates the effort to distinguish social machines from ostensibly similar systems, such as social computing, human computation, crowdsourcing and collective intelligence systems. The development of a taxonomic framework also marks an important step in our attempt to understand how social machines emerge and develop. Most importantly, the taxonomic framework establishes the dimensions according to which different instances of social machines can be compared to derive commonalities, as well as design and behavior patterns. This enables us to identify specific categories of social machines (taxa) that serve as the basis for classification efforts. It also enables us to analyze the overall space of design possibilities and identify areas that have been under-explored by research and development efforts. Finally, a taxonomic framework establishes the basis for future scientific efforts of both an analytic and synthetic nature: analytic efforts are driven by a need to understand why some parts of the design space are more populated than others, and synthetic efforts are driven by the need to explore parts of the design space that may afford opportunities for the creation of entirely novel kinds of social machines.

³See <http://sociam.org/>.

2 Social Machines: A Working Definition

Although there are a variety of views in the literature as to what actually constitutes a social machine, perhaps the most popular characterization is provided by Berners-Lee and Fischetti [3] in their book ‘Weaving the Web: The Original Design and Ultimate Destiny of the World Wide Web’:

Real life is and must be full of all kinds of social constraint—the very processes from which society arises. Computers can help if we use them to create abstract social machines on the Web: *processes in which the people do the creative work and the machine does the administration.* [our emphasis] (p. 172)

This characterization emphasizes the joint involvement of people and technology with respect to particular processes, and it also makes a distinction between the respective roles that people and machines play with regard to the process being undertaken; in particular, the contributions of the human participants should consist in some form of creative work, while the contributions of the machine components should consist in some form of administrative activity. Assuming that the notion of ‘creative work’ should be interpreted in terms of the generation of online content (e.g., uploading a photo or writing some text), then it seems that Berners-Lee and Fischetti’s understanding of social machines can be applied to many different kinds of Web-based systems. These include, for example, Wikipedia, Twitter, Facebook, YouTube,⁴ and Flickr.⁵ As is evidenced by the Web traffic system, Alexa,⁶ the sites that host these systems are among the most popular on the Web today.

Although Berners-Lee and Fischetti’s characterization can be used to support the identification of at least some social machines, it is far from clear that it applies to *all* social machines. One problem is that it is sometimes hard to discern what counts as a form of administrative and creative activity. Wikipedia bots,⁷ for example, engage in automated processes that are essential to the ways in which Wikipedia content is managed. In some instances, they use advanced machine learning techniques to perform tasks that not so long ago were exclusively tackled through manual work and human insight, for instance, to detect and remedy deliberate attempts to vandalize Wikipedia articles [33]. Such bot-related (i.e., machine-based) activities could be easily classified as ‘creative’. In other cases, we encounter human contributions that could be cast as somewhat administrative in nature. For example, the process of adding tags to Flickr photos plays an important role in terms of organizing the contents of the site, thereby making it easier for certain kinds of content to be accessed by the user community.

⁴<http://www.youtube.com/>.

⁵<http://www.flickr.com/>.

⁶<http://www.alexa.com/>.

⁷For an overview, see <http://en.wikipedia.org/wiki/Wikipedia:Bots/Status>, retrieved in December 2013.

Another problem with Berners-Lee and Fischetti's characterization is that it seems to overlook cases where the machine elements play an important role in the generation of online content or in enabling activities that are essential to it. PicBreeder,⁸ for example, is a system that supports the collaborative and interactive production of images using a mixture of evolutionary computation techniques and human agents [39, 40]. The role of the human agents here is to select the machine-generated images based on (e.g.) aesthetic criteria. These images are then published on the PicBreeder site and are accessible to other users who can use them as the starting point for their own interactive image generation activities. PicBreeder is thus a system in which the machine components arguably play an important role in terms of what appears online (it is, after all, the machine components that are generating the images). If we were to embrace the notion of social machines as systems in which it is the humans that are solely responsible for the creative work, then PicBreeder would seem to be a poor candidate as a social machine. And yet PicBreeder does seem to have many of the features that make it the legitimate target of attention for the social machine community: there is community engagement, issues of human-machine collaboration, the socially-distributed nature of particular tasks, and so on.

In view of these problems, we suspect that a definition that seeks to impose constraints on the precise nature of the contributions made by human and machine components with respect to the performance of a task (administrative, creative, or otherwise) is likely to be overly restrictive in terms of the identification of important and interesting social machine exemplars. More importantly, if we carry such notions forward into the design and development of social machines, we risk delivering systems in which the virtues of human-machine interaction with regard to creative (and administrative) processes are ignored. When it comes to creative processes, for example, we should recognize that some of our best creative accomplishments often come about as a result of our ability to engage and interact with our technological artifacts, and we should seek to exploit this in the context of our design and engineering efforts. A perspective that seeks to limit the kinds of roles that can be performed by human and machine elements, and which additionally seeks to impose a strict (and rather artificial) boundary on where particular processes need to take place, risks blinding us to many of the opportunities that the Web provides in terms of the transformation of traditional processes and the enhancement of both human and machine capabilities.

In response to the problems associated with Berners-Lee and Fischetti's characterization, Smart and Shadbolt [45] have proposed the following working definition of a social machine:

Social machines are Web-based socio-technical systems in which the human and technological elements play the role of participant machinery with respect to the mechanistic realization of system-level processes.

⁸<http://picbreeder.org/>.

This definition relaxes the constraint associated with the nature of the functional contribution made by human and machine elements, although it preserves the emphasis on social machines as bio-technologically hybrid systems (i.e., as systems that feature the incorporation of both people and machines). In particular, humans and machines are deemed to be *jointly* involved in the physical realization of processes: they are deemed to constitute part of the social machinery by which such processes are physically realized. This notion of human and machine elements serving as forms of participant machinery [8, p. 207] takes its inspiration from an approach to mind and cognition that sees extra-organismic resources as (on occasion) participating in the material realization of human mental states and processes—such resources are deemed to “form part of the very machinery by means of which mind and cognition are physically realized and hence form part of the local supervenience base for various mental states and processes” [8, p. 207]. A social machine is thus similar to what has been dubbed a ‘Web-extended mind’ [44] in the context of the Web Science literature.⁹ Essentially, we suggest that a social machine is an extended functional organization in which the explanation of certain system-level processes requires an account that adverts to the details of mechanisms that are distributed across both the biological (human) and the technological (conventional computing systems) realms.¹⁰ Such forms of ‘explanatory spread’ (see [58]) are sufficient for us to approach a social machine as a functionally-integrated system in spite of the heterogeneous nature of its material constitution. One of the crucial differences between the notion of a Web-extended mind and the notion of a social machine concerns the social aspect of the latter: the fact that it is multiple individuals (rather than a single individual) that participate in the realization of processes associated with the larger systemic organization. In addition, the kinds of processes enabled by the two scenarios are not co-extensive: Web-extended minds are concerned with cognitive processes; social machines, in contrast, are more general, referring to processes that may or may not be cognitive in nature.

Based on the above definition, a number of features of social machine systems are worth highlighting. One of these features concerns the fact that social machines are socio-technical systems—that is they involve the participation of human individuals and technological components. In many cases, we can expect the respective contributions of human and machine elements to draw on their distinctive capabilities and to complement one another with respect to the process that is being realized. It is the nature of this complementarity that underlies the interest in social machines as systems capable of a variety of advanced problem-solving capabilities

⁹The notion of a Web-extended mind draws its inspiration from work that goes under a variety of headings, such as ‘extended cognition’, ‘cognitive extension’ or ‘the extended mind’ [8, 9, 28]. Smart [44] defines a Web-extended mind as a system in which some of the informational and technological elements of the Web can be seen to constitute part of the material supervenience base for (at least some of) a human individual’s mental states and processes.

¹⁰The use of the term ‘mechanistic realization’ in the definition is intended to highlight the importance of this mechanistically-oriented explanatory account [59].

(see [18, 22]).¹¹ By virtue of their ability to factor in human and machine contributions, social machines are often able to extend the reach of both human and machine intelligence, supporting capabilities that less integrated systems might find difficult to accomplish. In the taxonomic framework introduced in Sect. 4 we will elaborate on this symbiosis with respect to the ways in which this integration is achieved in terms of task assignment mechanisms and the roles that each type of component plays in the overall system.

A second point that is worth emphasizing is that, for our purposes, social machines are cast as Web-based systems. Although we do not rule out the possibility of social machines that are independent of the Web,¹² we suggest that the properties of the Web make Web-based social machines a particularly important focus of social and scientific attention. One virtue of the Web, in this respect, is that it enables us to tap into the capabilities of human agents in a manner (and on a scale) that has never been seen before. The Web is a social technology that interfaces with a large proportion of humanity. By firmly embedding itself within a human social environment, the Web opens up a range of opportunities to incorporate human agents into episodes of machine-based processing. This makes Web-based social machines capable of supporting processes that would be difficult or impossible to realize in other kinds of social (or indeed socio-technical) context.

Thirdly, social machines are systems that consist of multiple (human) individuals. This aspect is crucial for understanding the capabilities of social machines and designing successful systems. By drawing on a large number of individuals, social

¹¹Similarly, it is the complementary nature of biological and non-biological resources (in terms of their contrasting representational and computational capabilities) that is often seen as lying at the root of the advanced forms of intelligence exhibited by extended cognitive systems. Sutton [49], for example, writes that “in extended cognitive systems, external states and processes need not mimic or replicate the formats, dynamics, or functions of inner states and processes. Rather, different components of the overall (enduring or temporary) system can play quite different roles and have different properties while coupling in collective and complementary contributions to flexible thinking and acting” (p. 194).

¹²Clocks may provide one example of a social machine that is independent of the Web. In their book, ‘Anti-Oedipus’, Gilles and Guattari [16] suggest that clocks are a form of ‘social machine’: “The same machine can be both technical and social, but only when viewed from different perspectives: for example, the clock as a technical machine for measuring uniform time, and as a social machine for reproducing canonic hours and for assuring order in the city” (p. 155). Interestingly, clocks have been seen as providing the technological impetus for the transformation of society. A number of theorists have emphasized the way in which clocks enable the large-scale scheduling and coordination of both individual and collective action, and the way in which the transition from fixed, centralized clock towers to portable wristwatches paved the way for new forms of social interaction and engagement [23]. The invention of portable time-keeping devices, argues Landes [23], made it possible to organize and synchronize activities in a way that had never been possible before, and on the back of this new capability there emerged a new social and economic era. The clock, in this case, can be seen as the technological element of a social machine in the sense that it is influencing social interaction via the delivery of machine-generated temporal representations. These representations serve to structure, sculpt and scaffold forms of social interaction and engagement that progressively shape the contours of the social, economic and cultural landscapes in which we live.

machines are able to accomplish tasks that require significant amounts of effort, for example, the decentralized analysis of large and complex bodies of scientific data (in Sect. 4 we will discuss the types of workflows that support this analysis at scale). In addition, social machines are able to exploit differences between individuals with respect to abilities, skills, insights, perspectives, knowledge, geographical location, experiences, group membership, social position, and so on. This may serve to improve the diversity and quality of the contributions that are made by the human community. Finally, social machines are also able to exploit the performance improvements that are often associated with collective inputs, for example, those associated with the Wisdom of Crowds phenomenon [48].

Fourthly, it follows from the above definition that processes are central to our understanding of what makes something a social machine: we discern a social machine when we encounter a process that demands a (mechanistically-oriented) explanatory account formulated in terms of the joint contributions of multiple individuals and Web-based technological components. It is important to note that we are not saying that social machines *are* processes, as would seem to be implied by the definition of social machines offered by Berners-Lee and Fischetti [3]. Rather, we are saying that social machines are the physical systems that perform, implement or realize such processes. This is an important distinction because the original definition (proposed by Berners-Lee and Fischetti) can result in a certain amount of confusion and conceptual indiscipline when it comes to discussions about social machines. Tinati and Carr [50] thus write that “any task that requires the co-constitutional involvement of humans and technologies is a form of social machine”. This characterization places appropriate emphasis on the importance of socio-technical engagement in the context of particular tasks, but it is a mistake to progress from this to the conclusion that the task itself is a form of social machine. Such conclusions, in our view, reflect a category error concerning the ontological status of social machines.

The centrality of processes to our understanding of social machines throws up a range of interesting issues and questions, some of which are out of the scope of the current chapter. One issue concerns the temporal nature of processes and the implication this has for the lifetime of a social machine. Processes may clearly be of relatively short-lived duration or they may be somewhat more enduring. Inasmuch as social machines exist for the duration of the processes with which they are associated, it would seem likely that social machines have a fair amount of variability with respect to their longevity. It should be possible to encounter social machines that persist for relatively long periods of time (as in the case of temporally-sustained, ongoing processes), as well as social machines whose existence is somewhat more fleeting and evanescent (as in the case of a social machine that supports social coordination in respect of a specific event—the organization of a birthday party, let’s say). Temporality plays a crucial role for several other properties of social machines captured by our taxonomic framework. For instance, the types of contributions made by human participants may change depending on their role in the system; also, the range of activities that are performed automatically might be

expanded by the availability of new algorithms (as was the case with the Wikipedia bots discussed earlier). Such temporal variability has implications for efforts that seek to observe and monitor social machines, such as the efforts associated with the Web Observatory initiative [10,51]. In particular, if we assume that persistent social machines are both easier to monitor and also generate the most data (on account of their temporally-enduring nature), then it becomes clear that we face the potential hazard of a sampling bias as part of our monitoring efforts. Equally important is how changes along one or several dimensions of our taxonomic framework (see Sect. 4) affect the frequency of the monitoring exercise and our ability to manage and derive insight from observational data. If our future scientific understanding of social machines is grounded on a limited subset of social machine exemplars (i.e., the long-lived ones), then it is unclear whether our understanding will ever be complete: the properties and dynamics of an entire class of perhaps socially- and cognitively-crucial systems will go unrecorded.

A second issue thrown up by the process-oriented nature of social machines concerns the nature and visibility of the goal that is being realized by the process. In some cases, the goal of the process that is being realized by the social machine will not be visible to the human participants in the system. In other cases, the goal may be visible to one or more of the human participants, perhaps because they are the ones responsible for assembling the social and technological elements into a functionally-integrated information processing ensemble. Importantly, it seems possible to discern some cases where a social machine may be created or emerge from a technological system that was originally designed or configured to perform a different function. A social machine that emerges in the context of a large-scale social networking service, for example, may be concerned with the modification of people's voting behavior (see [31]) or product consumption patterns and actually have very little to do with the formation and maintenance of social bonds. A second class of examples which exemplifies the varying degree of goal transparency/awareness can be found in the area of human computation. An important category of social machines are thus systems referred to as 'games with a purpose' (or GWAP) [55]. In such systems, human agents participate in a game, often interacting with each other, sharing scores and competing against friends from their social network. The inputs collected from the players, in this case, are used to improve the accuracy of computing algorithms; players are not necessarily aware of the actual goal of the game as it was conceived by the game designer, but their social interactions and game play result in useful training data that assists with the development of automated processes. This example also makes clear the ambivalent nature of such goals; one could distinguish among (sometimes overlapping) component-level and system-level goals, each equipped in some cases with a temporal element.

Finally, it is worth noting that social machines, in virtue of involving multiple individuals, are often concerned with processes that are relevant to the social interactions and relationships between individuals. Many of the processes in which human and machine elements participate may thus be glossed as 'social processes': they concern the structure and dynamics of a group of people. Such processes may be

many and varied. They include (but are not necessarily limited to) the coordination of collective action (e.g., implementations based on the Ushahidi platform); the pooling and distribution of resources (e.g., YouTube); the influencing of individual thoughts and actions (e.g., Twitter); the formation, maintenance, and dissolution of social relationships (e.g., Facebook); the collaborative creation of socially-shared assets (e.g., Wikipedia); and the social distribution of problem-solving processes (e.g., Galaxy Zoo). In general, the role of the machine or technological elements with respect to these processes is to constrain, control, coordinate or otherwise influence the social interactions between people (e.g., LinkedIn),¹³ or, alternatively, to govern the way in which individual human contributions are collectively factored into some other process (e.g., reCAPTCHA¹⁴). Typically, the influence exerted by the technological elements, in these cases, is mediated by some form of manipulation and processing of the informational inputs that are provided by human agents (this distinguishes social machines from systems which merely act as conduits for the communication of information between individuals). Alternatively, it is possible that the influence may be exerted through the provision of machine-generated representations; for example, system-generated cues play a role in governing the dynamics of person perception processes in the context of systems such as Facebook [52] and Twitter [57].

3 Examples of Social Machines

A broad range of Web-based systems have been considered as candidate social machines within the Web Science community. These include Facebook [18], mySpace [18],¹⁵ Twitter [17], YouTube [43], Ushahidi [41], Galaxy Zoo [17, 41], reCAPTCHA [30], Reddit [43],¹⁶ Wikipedia [17, 18, 41], Amazon's Mechanical Turk [43],¹⁷ and the Web itself [17].¹⁸ As should be clear from this list, social machines are a pretty heterogeneous bunch of systems. For one thing, they seem to occupy a variety of functionally-diverse niches within the ecology of the Social Web. Extant social machines thus include social networking systems (Facebook, mySpace, Twitter), microblogging services (Twitter), video/photo sharing systems (YouTube), citizen science projects (Galaxy Zoo), social news sites (Reddit), collaborative content editing sites (Wikipedia), frameworks for the creation of collaborative systems (Ushahidi) and systems that enable human contributions to

¹³<https://www.linkedin.com/>.

¹⁴<http://www.google.com/recaptcha>.

¹⁵<https://myspace.com/>.

¹⁶<http://www.reddit.com/>.

¹⁷<https://www.mturk.com/mturk/>.

¹⁸The Web site of the SOCIAM research project lists a large number of additional examples of social machines—see <http://www.sociam.org/social-machines>.

be productively exploited in the context of automated processes (reCAPTCHA) or more traditional production processes (Mechanical Turk). This diversity has implications for the kind of features that we rely on to discriminate between social machines (see Sect. 4), and it also has implications for the types of social machines that we are able to recognize. The aforementioned list of social machine exemplars also highlights a number of areas of confusion when it comes to an understanding of the social machine concept. Armed with the working definition from Sect. 2, we are now in a position to address these areas of confusion (see also Fig. 1).

The first thing to note is that it is very common for people to refer to specific technologies when they talk about social machines. In many cases, therefore, when people identify a given social machine instance they point to a platform such as Facebook, Twitter or Ushahidi. Figure 1 refers to these as ‘frameworks’ and ‘services’, thus emphasizing the key role these socially-active environments play in the development and emergence of a wide range of special-purpose social machines targeting less general audiences. It is important to be clear that when we talk about social machines we are talking about a *socio-technical* system (as opposed to a purely technological system) that is actively engaged in the realization of a particular process [43]. Thus, when we say that Facebook is a social machine, what we mean is that it is the social networking platform (that we typically identify as Facebook) plus the human participants (the social environment) that constitutes the social machine. Any reference to a social machine as being constituted *solely* by the technological system (or subsystem) is, in our view, incorrect. It is for this reason that it is probably a mistake to refer to the human components of a social machine as the ‘users’ or as forming part of the ‘user base’ of the social machine.¹⁹ Such terms imply that the social machine is something separate from the human participants: it conjures up an image of social machines as things that are independent of the human communities with which they are associated, and it encourages us to place undue emphasis on the technological aspects of the system. As should be clear from the definition presented in Sect. 2, social machines should be properly conceptualized as socio-technologically integrated systems in which the human ‘users’ are an intrinsic part of the larger, biotechnologically-hybrid system. This does not, of course, undermine the importance of the technological aspects as a source of scientific interest and a focus of engineering attention. Even in cases where all forms of human participation are absent, we can still recognize a technological system as something that is apt to participate in the formation of a social machine (or a multiplicity of such machines), and treat it as a legitimate target of scientific enquiry. The fact that an aircraft carrier is not, by itself, a socio-technical system does not mean that such vessels are not of considerable interest to naval engineers, even in situations where it is clearly obvious that the processes that the vessel is designed to support could only be realized once the human crew is onboard and certain forms of socio-technical entanglement are established.

¹⁹We are grateful to Ségolène Tarte (University of Oxford) for pointing this out.




Fig. 1 Distinction between technology and social machines as part of a broader ecosystem [43]

A second issue arising from the aforementioned list of social machine exemplars concerns the distinction between systems that serve mainly as frameworks for the creation of social machines and systems that actually function as social machines (see the framework tier of Fig. 1). Ushahidi, for example, has been used to develop a number of systems that support information collection, visualization and interactive mapping, as in the Ushahidi-based system that supported humanitarian relief efforts in the aftermath of the 2010 Haitian earthquake [60]. An analogy here is with the MediaWiki system, which has supported a wide range of wiki-based collaborative content creation projects (e.g., Wikipedia, Wiktionary, Wikidata and Wikispecies). Wikipedia and Ushahidi are not, therefore, instances of the same class of objects, as might be implied by the above list of social machines. Instead, Ushahidi and MediaWiki should be seen as frameworks that support the creation of specific systems, such as Wikipedia and the Haitian implementation of Ushahidi. Another example of a framework that can be used to support the creation of social machines is Diaspora.²⁰ It can be used to create social networking services for specific communities of users. Although such frameworks should be distinguished from actual instances of social machines, they are clearly relevant to the project of characterizing and understanding social machines. For one thing, such systems serve as the template for a range of social machines that may possess similar or identical characteristics, and their design greatly influences the way in which a social machine functions and evolves.

A third point of interest concerning the aforementioned list of social machines concerns the way in which some social machines emerge in the context of other systems, which themselves may or may not be regarded as social machines.²¹ O'Hara [30], for instance, talks about the use of Facebook to organize a birthday party. In this case, it is the specific use of Facebook to accomplish a particular task (i.e., organize a birthday party) that counts as a social machine rather than (perhaps) a more liberal perspective that sees Facebook itself as a social machine: Facebook is, in this case, merely serving as a form of technological scaffolding that supports the creation of a multiplicity of (probably) short-lived social machines. A similar claim could be made with respect to the relationship between the Web and social machines. Thus, although the Web has been regarded as a social machine [17], perhaps it is more appropriate to see the Web as the technological matrix that gives rise to a variety of social machine systems and in which all such systems are ultimately embedded. Contrary to this interpretation, however, we might argue that nothing in the definition of a social machine—either the original characterization [3] or the more recent definition [45]—would seem to rule out the possibility of

²⁰<https://diasporafoundation.org/>.

²¹This corresponds to the tier termed 'Causes/Groups' in Fig. 1, which builds on a selection of Web-based systems that, through their large-scale user bases and general character, have reached a level of popularity that turns them into frameworks for the development of more special purpose social machines.

either the Web or Facebook counting as social machines. In addition, the possibility of a social machine emerging from the material matrix associated with some other system does not rule out the possibility that the other system is in fact a social machine: it may just be that the material elements of one social machine (i.e., its human and technological components) are simply recruited to form a social machine that is involved in a different process.²² We suggest that we tend to discern a social machine when we can identify a socio-technical system that is involved in the realization of processes associated with the performance of a particular task. With this in mind, we might feel inclined to see a distinct social machine (one that draws on the technological fabric of Facebook, let's say) whenever we see particular tasks being performed (e.g., organizing a birthday party). However, in many cases, the larger system is also involved in the performance of particular tasks. Thus, in the case of Facebook, we might say that the system is (broadly) engaged in the realization of (the more temporally-protracted) process of social relationship management (i.e., the creation, maintenance, and dissolution of social networks). Inasmuch as we see this process as one in which the technological elements of the Facebook system are playing an explanatorily significant role, then we see no problem with a perspective that views Facebook as part of a functionally-integrated system (i.e., a social machine). Obviously, this does not rule out the possibility that the material elements associated with this system could be involved in a multiplicity of other, perhaps more short-lived, processes.

From an engineering point of view, the realization of such ecosystems depends on technologies, services and generic platforms that not only provide specific functionality—depending on the kind of social process supported by the social machine, this could be anything from communication and coordination of joint efforts to collaborative content generation, knowledge sharing, and decision making—but also promote principles, values, and ideas that match the expectations and motivation of the human participants. In particular, due to the very nature of a social machine and its ecosystem, it is essential that the technologies used to realize it are equipped with the means to tackle scale, decentralization, and concurrent access and processing. As content is created and shared in a distributed fashion, the social machine must be able to establish and associate trust or at least accountability in the ways every component of the system, biological or technological, operates and interacts with the rest of the ecosystem. We will follow up on these aspects in Sect. 4 where we discuss the social machine taxonomic framework.

²²It is also possible to imagine one or more social machines being 'incorporated' into a larger social machine. In the same way, perhaps, as the neurological subsystems associated with memory, attention and perception merge to form part of the integrated mechanistic substrate that realizes more 'macrocognitive' functions such as sensemaking (see [21]).

4 Characterization of Social Machines

As part of the attempt to understand social machines, it is useful to develop a taxonomic framework that can be used to describe and classify social machine instances. Following Nickerson et al. [29], we define a taxonomy T as a set of n dimensions D_i ($i = 1, \dots, n$) each consisting of k_i ($k_i \geq 2$) mutually exclusive and collectively exhaustive characteristics C_{ij} ($j = 1, \dots, k_i$) such that each object (i.e., social machine) under consideration has one and only one C_{ij} for each D_i . We have adopted this definition for our own taxonomic framework. Our approach to taxonomy development is also based on the approach advocated by Nickerson et al. [29], which has its roots in the social sciences (see [2]). The approach consists of three stages (see Fig. 2):

1. **Empirical-to-Deductive Stage:** This stage involves the initial examination of a subset of objects (social machine instances in our case) and the identification of their distinguishing features. As will be clear from the subsequent discussion, we rely on specific techniques in order to support this process. The output of this stage is an initial set of dimensions.
2. **Deductive-to-Empirical Stage:** This stage entails the conceptualization of new characteristics and dimensions. The dimensions elicited in the empirical-to-deductive stage are progressively refined and enriched during this stage.
3. **Taxonomy Application Stage:** This stage involves the use of the taxonomy to identify and characterize new objects. The taxonomy may also be used to inform the design of new objects.




Fig. 2 The three stage approach to taxonomy development that was adopted in the current study (see Nickerson et al. [29] for more details)

As part of the empirical-to-deductive stage, we constructed an initial social machine taxonomy that included a set of dimensions but lacked any characteristics. This work is summarized below and reported in more detail in Shadbolt et al. [43]. In the current chapter, we focus on the deductive-to-empirical stage and present a more complete taxonomy featuring a revised set of dimensions and a complete set of characteristics. Although this taxonomy is subject to further refinement, it is suitable for use within the final taxonomy application stage of the taxonomy development process outlined above. In particular, we have compared the results of the taxonomy development effort with similar efforts that have been made in related areas (see Sect. 5). As part of our future work, we will test the completeness, correctness, and comprehensibility [56] of the taxonomy in experiments in which a new set of social machines will be classified by framework users. We will ask the participants to assess the quality of the framework along these general dimensions, and measure inter-annotator agreement to learn about the usefulness of the classifications produced.

We now turn to a description of the current version of the taxonomy. We will first present the approach taken to elicit information about social machine dimensions from human subjects and then give a summary of the empirical and conceptual work undertaken to define the taxonomy dimensions and their associated characteristics.

4.1 Eliciting Social Machine Dimensions

As illustrated by the examples surveyed in Sect. 3, social machines come in a variety of shapes and sizes. A system such as Facebook, which is concerned with (among other things) the formation and maintenance of social relationships, is clearly different from a system such as Mechanical Turk, which offers a crowdsourced labor market for simple data collection and processing tasks, and both of these are different from Zooniverse, which supports a form of socially-distributed problem solving in the natural and life sciences. Based on the working definition introduced in Sect. 2, we can anticipate a number of ways in which social machines built around these kinds of technological systems might differ. These include differences in the nature of the processes being realized; the kinds of contributions or computations made by human and machine components; the relative balance of processing effort among these components; and the ways that individual contributions are combined in the course of process execution. These, however, are just a few examples of the dimensions that could be used to differentiate between social machines. Other dimensions might be less obvious based on a cursory analysis of a limited subset of what are arguably the most well-known social machines that currently exist. Furthermore, even when larger samples of social machines are surveyed, the task of eliciting a more-or-less complete set of dimensions is not straightforward. People may find it difficult to discern differences between social machines, or find it difficult to communicate their conceptual understanding of these systems in a structured and coherent way, even when they are using these systems on a regular basis.

One way of dealing with the difficulty of eliciting dimensions is to rely on a range of techniques known as knowledge elicitation techniques [42]. These techniques are used as part of knowledge engineering initiatives in order to create the conditions under which domain experts are best able to communicate the knowledge associated with their expertise. Although a broad range of techniques are available, the ones that tend to be most suited for the elicitation of information about the dimensions along which a set of common objects vary are sorting and rating techniques. These include the repertory grid technique, which has its roots in the psychology of personality [13, 19, 20]. The repertory grid technique is useful because it can be used to support the elicitation of knowledge that is largely tacit in nature, i.e., difficult for an individual to verbalize. In addition, the data that is obtained as part of the technique can be subjected to various forms of statistical analysis in order to obtain insight into the structural organization of domain-relevant concepts.

In a repertory grid exercise subjects are presented with a range of objects, referred to as 'elements', and asked to choose three, such that two are similar and different from the third. This is known as the method of triadic elicitation (e.g., [6]). In order to demonstrate the technique, imagine a subject is presented with the following set of social machines²³: Facebook, Twitter, YouTube, Wikipedia, reCAPTCHA, Galazy Zoo, Flickr, mySpace, LinkedIn,²⁴ and Planet Hunters.²⁵ The participant might choose Twitter and Facebook as the two similar elements and Galaxy Zoo as different from the other two. The subject is then asked to provide the reason for differentiating these elements, and this dimension is known as a 'construct'. Each construct is assigned a name as are the two poles that represent the opposite ends of the construct. In our example, 'size of the user community' might be a suitable construct that differentiates between the selected elements, with 'small' and 'large' serving as the two poles of the construct. Once a construct has been elicited, all the elements can be rated with respect to the construct, with the ratings reflecting the extent to which the subject sees an element as falling to one or other of the construct poles. The process of triadic elicitation is continued with different triads of elements until the subject can think of no further discriminating constructs. At this point we have a matrix of similarity ratings that can be analyzed using techniques such as cluster analysis. This provides us with a dendrogram that can reveal conceptually-significant categories of social machines, and it can also shed light on the relationships that exist between the constructs.

In order to test the applicability of the repertory grid technique to the social machine taxonomy development effort, we first undertook a knowledge elicitation session with a computer science researcher from our laboratory. We presented them with the ten social machines listed above and engaged them in a process of triadic elicitation, in each case asking them to select and discriminate between

²³In fact, as we mentioned in Sect. 3, the technological subsystem is only considered to be one part of the social machine; the human participants are also deemed to be part of the social machine.

²⁴<https://www.linkedin.com/>.

²⁵<http://www.planethunters.org/>.

elements. The following set of constructs were elicited as a result of the application of this technique (elements that are representative of the poles of each construct are presented in square brackets):

1. **Size of the User Community:** the number of users that participate in the system, either as contributors or consumers [small: Galaxy Zoo; large: Facebook].
2. **Extent of User Contribution:** the proportion of users that actually contribute content as opposed to users that merely consume content [small: Wikipedia; large: Galaxy Zoo].
3. **Sociality:** the extent to which the system supports social interaction [low: reCAPTCHA; large: Facebook].
4. **Visibility of Individual User Contributions:** the extent to which individual user contributions are visible to the entire user base of the system [low: reCAPTCHA; high: YouTube].
5. **Inter-dependence of User Contributions:** the extent to which the user contributions are independent of one another with respect to the task being performed by the system [low: Twitter; high: Wikipedia].
6. **Focused on a Single Task:** the extent to which the social machine is focused on a single task as opposed to supporting multiple kinds of tasks [single task: Galaxy Zoo; multiple tasks: Facebook].
7. **Anonymity of Human Users:** the extent to which the system requires users to provide personal information about themselves to other users [low anonymity: Facebook; high anonymity: reCAPTCHA].
8. **Heterotelic vs. Autotelic Usage:** the extent to which the use of the system is motivated by instrumental or professional (heterotelic) concerns as opposed to enjoyment and pleasure (autotelic) [heterotelic: LinkedIn; autotelic: YouTube].
9. **Requires the Aggregation of User Contributions:** the extent to which user contributions need to be aggregated in order for the social machine to perform its primary function [low: Twitter; high: Wikipedia].
10. **Diversity of User Contributions:** the degree of differentiation with respect to user contributions. For example, users may be engaged in a single task (e.g., galaxy classification) or multiple tasks (e.g., uploading, rating, and tagging content) [low: reCAPTCHA; high: YouTube].
11. **Salience of Social Network:** the relative salience or visibility of the social network within the system [low: reCAPTCHA; high: Facebook].

The rating matrix and results of the cluster analysis are presented in Fig. 3 (these results were obtained using the WebGrid 5 system).²⁶ The ten social machines that were the focus of the repertory grid (e.g., reCAPTCHA, Galaxy Zoo, Planet Hunters, etc.) are listed at the base of the rating matrix, and the labels used to describe the poles of each construct are listed on either side of the matrix (e.g., ‘heterotelic use’ vs. ‘autotelic use’). The numbers that make up the body of the matrix are the ratings made by the user for each of the social machines, with lower

²⁶See <http://gigi.cpsc.ucalgary.ca/>.

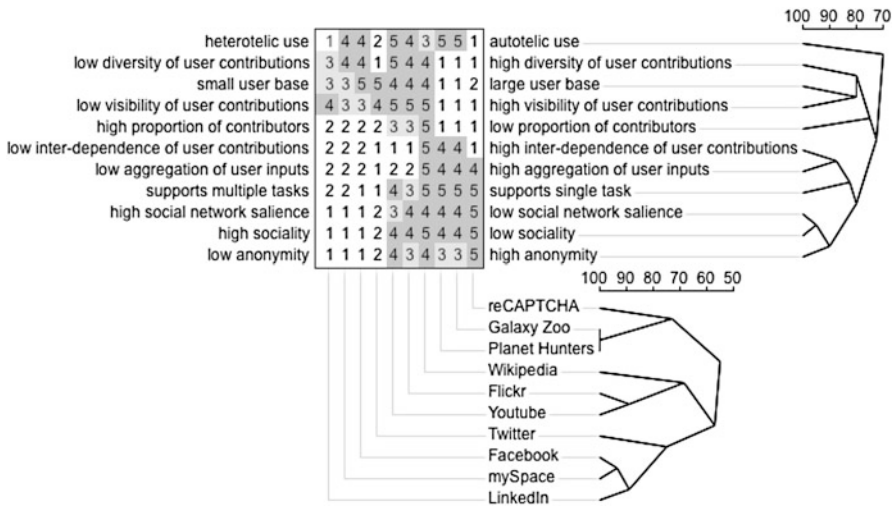


Fig. 3 The results of the repertory grid technique applied to the domain of social machines [see text for a description of the rating matrix and dendograms for both the constructs (*top dendogram*) and the social machine elements (*bottom dendogram*)]

ratings reflecting a bias towards the pole on the left hand side of the matrix. A score of ‘1’ in the case of the ‘Heterotelic vs. Autotelic Usage’ construct thus indicates that the social machine is used for heterotelic purposes, whereas a score of ‘5’ indicates that the social machine is used for autotelic purposes. With respect to Fig. 3, we can therefore see that LinkedIn and reCAPTCHA are both examples of social machines that are used predominantly for heterotelic purposes (they both have a low rating), whereas Galaxy Zoo and YouTube are both used predominantly for autotelic purposes (they both have a high rating).

The first thing to note from the dendogram associated with the social machines is that the Planet Hunters and Galaxy Zoo systems emerge as identical systems in this analysis—there are no constructs that differentiate between these two elements. This presumably stems from the fact that both systems form part of the Zooniverse²⁷ collection of citizen science projects and both are concerned with the analysis of astronomical data. This result could be used to elicit additional differentiating constructs in situations where the subject did in fact believe there to be differences between the two systems. Another feature to note from the dendogram is that Facebook, mySpace and LinkedIn all seem to form a distinct cluster. We can ask our subject to attempt to say something about this clustering, perhaps supplying a label to identify a class or category of systems. The response, in this case, could be that the systems are all examples of ‘social networking systems’. Another category of social machines seems to emerge based on the similarity of YouTube and Flickr.

²⁷<https://www.zooniverse.org/>.

In this case, we might say that these systems are both examples of ‘media sharing systems’.

In addition to the dendrogram associated with social machines, Fig. 3 also shows the dendrogram associated with the constructs. Here we can detect a number of correlations between the similarity scores, and these may reflect interesting contingencies between the features of social machines. For example, systems that exhibit low sociality also tend to be systems in which the social network has low salience. In addition, such systems are also ones that feature high levels of anonymity with respect to user contributions. As one might expect, systems that aim to support social interaction tend to require the disclosure of personal information—such disclosures are, in fact, likely to be a prerequisite for the development of relational intimacy. Another correlation emerges between the inter-dependence of user contributions and the tendency to aggregate user inputs. Again, not surprisingly, systems that feature high levels of interdependence between tasks also tend to be systems that engage in some form of aggregation of the user inputs. As part of our future work, we plan to collect a much larger collection of classifications in order to support the quantitative analysis of these sorts of correlations.

As should be clear from this example, the repertory grid technique can serve as an effective means of eliciting information about the features of social machines. It can also provide insight into the structure of the conceptual landscape associated with social machine systems. In particular, as more and more objects are surveyed, one can use cluster analysis to reveal interesting groupings that may serve as the basis for hierarchically-organized conceptual categories (i.e., taxa within the taxonomic framework). The results of the analysis can also serve as the basis for more focused knowledge elicitation sessions. For example, with respect to the above analysis, we could attempt to differentiate between the Planet Hunter and Galaxy Zoo systems, or we could exploit the ability to identify conceptual categories as a means of identifying additional social machines (e.g., systems that are members of the categories ‘social networking system’ and ‘media sharing system’).

4.2 A Social Machine Taxonomy

The analysis of the repertory grid described in the previous section provides some insight into the dimensions associated with social machines.²⁸ However, in order to expand the range of constructs elicited, it is necessary to draw on the perspectives of multiple individuals with respect to different subsets of social machines. For this reason, we completed an extended study involving ten computer science researchers

²⁸The constructs identified in the context of the repertory grid exercise ultimately drive the generation of dimensions associated with the taxonomic framework. A construct such as ‘Heterotelic vs. Autotelic Usage’ (see Sect. 4.1), for example, is ultimately used as the basis for the ‘Motivation Type’ and ‘Form of Motivation’ dimensions listed in Table 2.

from our laboratory [43]. The motivation for using computer science researchers, in this case, relates to the requirements of the repertory grid technique. In particular, the repertory grid technique requires subjects to be familiar enough with the elements being investigated in order for them to make meaningful comparisons and identify distinguishing features. Given that the computer science researchers in our laboratory are currently involved in the analysis of a broad array of social machines, it made sense to draw on their experience in the context of this particular exercise.

After each subject had completed the repertory grid analysis with their self-selected elements, the result was a set of ten repertory grids containing a combined total of 117 different constructs and 56 unique social machine instances. This marked the completion of the empirical-to-deductive phase of taxonomy development. We subsequently reviewed these constructs to identify closely related ones, grouped the resulting list into broader categories, and refined the taxonomy based on insights gained from a review of the relevant Social Web literature.

The results of this second, deductive-to-empirical stage of taxonomy development are presented in Tables 2, 3, 4, 5, and 6 (see appendix). We identified a total of 33 dimensions, which were organized into five categories. The categories relate to the tasks that are being performed by the social machine (or the processes being realized by the social machine), the (human–human, human–machine, and machine–machine) interaction mechanisms by which the social machine operates, the ways in which quality and performance are assessed, the motivational factors and incentive mechanisms that govern user participation in the system, and the technologies used to implement the technical grounding of the system. Across the 33 dimensions, we identified a total of 106 distinct characteristics.

4.3 *The Social Machine Morphospace*

The dimensions revealed by our analysis constitute the set of dimensions along which all social machines (extant or otherwise) can be deemed to vary. These dimensions can be used to define the axes of a multi-dimensional design space for social machines. This design space constitutes the universe within which all theoretically possible social machines are located, with the location of each social machine dictated by the particular combination of characteristics it possesses. Given the similarity of this design space to the notion of a ‘morphospace’ in the biological literature [36, 53], we refer to the design space (or universe of social machine possibilities) as the ‘social machine morphospace’. As with its biological counterpart, the social machine morphospace aims to chart the space of social machine possibilities with respect to a set of common features (dimensions) along which all social machines vary.

One advantage of the taxonomic framework is that it allows us to assess how much of the design space for social machines has been explored by current development efforts (obviously, given the size of the morphospace, it is likely that

this space will be sparsely populated). Regions within the space that are devoid of social machines may represent unexplored regions that provide fertile ground for the creation of novel systems. Alternatively, it may be that such regions are barren for a good reason: perhaps the design candidates that occupy this region are impractical or impossible to implement. In summary, the value of the social machine morphospace is that it provides a view as to the total space of design possibilities for social machines, and it indicates the regions of this space that have been unexplored by current development efforts. Not only does this shed light on the possible nature of future social machines, it may also help us to identify the specific combination of characteristics that determine whether a particular social machine fails or flourishes within the (current) socio-technical ecology of the Web.

5 Related Work

Given the value of taxonomies in advancing our understanding of the conceptual landscape associated with a domain, it is no surprise to discover that taxonomies have been developed for a range of systems appearing in the context of the Social Web. This includes, most notably, crowdsourcing [11, 14] and human computation systems [34], although similar attempts at characterization have been made in respect of social computing [1] and collective intelligence systems [25]. While none of the concepts associated with these systems are synonymous with the notion of social machines (see [43]), there are clear relationships between these various concepts. Instances of at least the technological components of social machines (e.g., Facebook, Wikipedia, Galaxy Zoo, etc.) are sometimes presented as instances of other kinds of systems, and this suggests that some of the dimensions associated with the social machine taxonomy may also surface in the context of other taxonomies. In order to evaluate this, we systematically compared the dimensions listed in Tables 2, 3, 4, 5, and 6 (see appendix) with those appearing in other studies [1, 11, 14, 25, 34]. The results of this analysis are presented in Table 1 (see appendix). As can be seen from this table, a number of social machine dimensions have at least some partial mapping to the dimensions identified in other studies.²⁹ This is particularly noticeable in the case of human computation and crowdsourcing systems (although this may simply reflect the greater attention that has been afforded to these systems in the context of previous taxonomy development efforts) [11, 14, 34].

²⁹Note that although two dimensions may be similar, they are only regarded as identical if the set of characteristics associated with the dimensions is the same in each case. In the absence of shared characteristics, a dimension mapping is regarded as ‘partial’.

6 Future Work

The definition of social machines presented in Sect. 2 and the taxonomic framework presented in Sect. 4 form part of an integrated attempt to develop a conceptual foundation for social machine research. It should be clear, however, that much more work needs to be done to make progress in this area. In terms of our conceptual understanding of social machines, for example, a range of perspectives exist concerning the nature of social machines. The definition of social machines that we have adopted here (and also in [45]) emphasizes the role of human and technological elements in the joint realization of processes. We might refer to this as the ‘socio-technical perspective’ of social machines. Such a perspective is, however, only one among many alternative perspectives that could be countenanced. While our definition is largely consistent with the views expressed by others in the Web Science community,³⁰ there are a number of competing perspectives available, and these need to be given closer scrutiny. An alternative concept, for example, tends to see social machines as socio-computational systems. According to this view, social machines are socially-extended computational systems in which some aspects of the computational process are delegated to multiple human individuals. Perhaps unsurprisingly, this kind of view tends to emerge in discussions of what has been dubbed ‘the social computer’ [38]. While there is clearly a certain amount of common ground between the ‘socio-computational perspective’ and the ‘socio-technical perspective’ (e.g., both regard social machines as systems that implement certain types of processes), there is a significant difference in terms of the scope of the conceptualizations entertained by each perspective. In particular, the socio-computational perspective seems committed to the view that social machines exist as a specialized form of human computation system [24]. We suggest that this contributes to an unproductive narrowing of the scope of social machine research efforts: it limits our scientific remit to a subset of Web-based systems whose constituent processes can be properly described as ‘computational’ in nature. In addition, by casting social machines as a specialized form of human computation system, we allow the scientific effort associated with the study of social machines to be too easily subsumed within an existing, and well established, field of scientific enquiry. In our view, the term ‘social machine’ is best reserved for a class of systems whose most important distinguishing feature is the manner in which system-level processes are realized. This is preferable to a perspective that focuses on issues of whether the process in question is or is not computational in nature. The crucial difference between the two perspectives is highlighted by the emphasis the

³⁰Such consistency is evidenced by the way social machines are described in a number of papers. We thus encounter descriptions of social machines as “purposefully designed sociotechnical system[s] comprising machines and people” [10], as systems in which “the human and digital parts. . . [form] a machine in which the two aspects are seamlessly interwoven” [43], and as systems that involve “the co-constitutional involvement of humans and technologies” [50].

socio-technical perspective places on the *way* in which a process is realized (i.e., the details of its mechanistic realization); the issue of whether or not the process in question can be characterized in computational terms is largely irrelevant.

A further focus area for conceptual analytic efforts is to distinguish between the notion of a social machine from a variety of ostensibly similar notions. These include crowdsourcing [7, 11], human computation [34], collective intelligence [25], social computing [32], the global brain [4], the social computer [38] and the social operating system [35]. It has been suggested that the social machine concept is similar to but not synonymous with (at least some of) these other concepts [43]. Additional work is required, however, to elucidate the exact nature of the relationships between the concepts. Furthermore, it will be important to ascertain the degree of overlap in the extensional projections of the concepts expressed by these terms.

As a means of furthering the effort to improve our conceptual understanding of social machines, we may be able to extend the methodological approach that was adopted in the case of the taxonomic framework; i.e., we may be able to make use of a range of knowledge elicitation techniques. Aside from the repertory grid technique (described above), a number of other knowledge elicitation techniques are available, and these could be useful in terms of exploring the social machine conceptual landscape. These include laddering techniques (useful for eliciting hierarchically-organized classes of social machines), concept sorting techniques (useful for identifying the features of social machines) and concept mapping techniques (useful for identifying the relationships between social machines) (see [42]). As with other applications of knowledge elicitation techniques, the results of these studies could serve as the basis for ontology development efforts. Such ontologies could then be used to provide machine-readable characterizations of specific social machine instances.

Regarding the effort to develop a taxonomic framework for social machines, a number of further steps need to be undertaken. Following the methodology advocated by Nickerson et al. [29], our work to date has focused on the empirical-to-deductive and deductive-to-empirical stages. The aim of the third stage—taxonomy application—is to use the taxonomic framework to identify and characterize additional instances of social machines. By situating these instances within the social machine morphospace, we will be able to chart the location of unexplored or under-explored regions of the design space. Of course, given the number of dimensions and the number of potential social machines that may emerge in the context of the current and future Web (recall that the study by Shadbolt et al. [43] yielded an initial sample of 56 social machines), the task of taxonomy application is likely to be something of a laborious undertaking (at least when seen from the perspective of a single individual). Clearly, one strategy for dealing with these sorts of tasks is to draw on the (socially-situated) processing resources made available by the technological infrastructure of the World Wide Web. This is precisely the strategy taken by social machines and other kinds of systems within the context of the Social Web. An interesting possibility, therefore, is to engineer a social machine

to expedite the process of taxonomy application.³¹ One specific idea that is currently under development is to build a microtask environment, including specific game elements, in which participants are asked to provide answers to atomic challenges that rate and compare a pair of social machine instances according to a dimension in our framework. Such systems may serve as a useful adjunct to ongoing initiatives, such as the Web Observatory initiative [51], which seek to observe the behavior of social machines within the ecological environment of the Web [10].

The use of the taxonomic framework to characterize new social machine instances is also a useful way of validating the framework. In particular, the attempt to characterize novel social machines enables us to answer questions concerning the generality (e.g., can we specify characteristics for all social machines?), accuracy (e.g., are the characteristics associated with a particular dimension mutually exclusive for any given social machine, or can a social machine have multiple characteristics on the same dimension?) and reliability (e.g., is the same system characterized in the same way by multiple users?) of the framework.³² We may, of course, discover at this stage that some putative social machines cannot be accommodated within the taxonomic framework. This may point to an inadequacy of the framework, or (more positively) it may indicate that the system in question is not, in fact, a social machine. In other words, the taxonomic framework could (ultimately) serve as a useful means of identifying bona fide members of the class of social machines. There are a number of different methodologies in the knowledge engineering literature which describe the steps to be followed in order to carry out the validation, and means to measure and analyze different validation criteria (see, for instance, [56]).

The use of the taxonomic framework to identify and characterize social machines yields a range of benefits. Firstly, by situating social machines within the social machine morphospace, we are able to determine the degree of clustering within the design space. We are able to answer questions concerning the extent to which existing systems are clustered together (like stars within a galaxy) or whether they are more-or-less randomly distributed across the void. This helps to determine whether current design and engineering efforts are focused on particular regions of the design space. Secondly, the population of the morphospace enables us to imagine as yet unrealized forms of social machines. By supporting our ability to focus on previously unexplored regions of the morphospace, the taxonomic framework is functioning as a ‘cognitive scaffold’ for our imaginative efforts. Such efforts

³¹Note that in the light of our definition, the ‘engineering’ of a social machine entails more than just software development and deployment; it also includes the assembly of mechanisms that enable and encourage user engagement.

³²The reliability of the framework is indicated by inter-rater reliability metrics. Poor measures of inter-rater reliability may indicate that some dimensions are more difficult to interpret, understand or discern than others. This may call for the dimension to be refined or removed from the framework.

may feed into the design and development of new kinds of social machines. Thirdly, we can use the body of data associated with the characterization of social machines in order to support efforts aimed at identifying categories or classes of social machines (using quantitative methods). We have alluded to a number of these categories earlier in the paper. For example, as a result of the repertory grid analysis described in Sect. 4.1, we made reference to ‘social networking systems’ and ‘media sharing systems’. Other classes of social machine focus on certain vertical sectors, for instance ‘crime social machines’ [12] and ‘health social machines’ [54]. Clearly, the effort to develop a hierarchically-organized set of social machine classes is an important focus area for future work,³³ and it could feed into the aforementioned effort to develop a social machine ontology. Finally, the application of the taxonomic framework yields a body of data that can be used to assess the relationship between particular combinations of characteristics and a range of interesting properties relating to (e.g.) the performance profile of the system and the size of its user community. These kinds of properties tend to be ones that determine how ‘successful’ a social machine is (e.g., whether it is able to achieve the goals its designers originally intended it to achieve), and thus the collection of correlational data is potentially useful in terms of guiding the design and development of new machines, as well as configuring existing ones. It should also be noted that the dimensions associated with the taxonomic framework can serve as independent variables in the context of experimental efforts intended to elucidate the relationship between particular characteristics and performance outcomes. Such variables are particularly useful in the case of cognitive social simulation studies where large numbers of cognitively-sophisticated agents can be used to shed light on the complex interactions between factors spread across the technological, informational, social and cognitive domains [47]. They can also offer a useful empirical grounding for system designers and inform the engineering and evolution of existing systems. One element of our future work in this area aims to investigate the dynamics of social machines using a combination of multi-agent simulation and cognitive modeling techniques (see [37]). From an engineering and HCI perspective, we will analyze the data collected through the application of the framework to derive best practices and guidelines for system design, which might also prove useful for ongoing initiatives such as the Web Observatory initiative.

³³The process of identifying categories or classes of social machines is supported by the use of statistical methods that are applied to the social machine morphospace. Cluster analytic techniques are typically used to support these analyses (see Geiger et al. [14] for an example of such techniques applied to crowdsourcing systems).

7 Conclusion

The recent growth and influence of the Social Web has led to an intensification of research efforts to understand the nature and dynamics of Web-based socio-technical systems. As part of these efforts, the term ‘social machine’ has emerged to help focus attention on a specific class of systems and to help delineate a range of theoretical, empirical and engineering issues. Although there is still no consensus regarding the precise semantics of the term ‘social machine’, we suggest that the notion of a social machine can best be understood in terms of particular processes (i.e., ones in which our explanatory accounts need to advert to the details of social participation and bio-technological coupling). We thus endorse the following definition of social machines:

Social machines are Web-based socio-technical systems in which the human and technological elements play the role of participant machinery with respect to the mechanistic realization of system-level processes.

As part of the effort to improve our conceptual understanding of social machines, we have attempted to construct a taxonomic framework. This framework draws on previous work that relied on the use of knowledge elicitation techniques to capture information about the various dimensions along which extant social machines can be deemed to vary [43]. We have extended this initial work by refining the set of elicited dimensions and also identified discrete characteristics for each of the dimensions. The result is a taxonomic framework consisting of a total of 33 dimensions and 106 characteristics. This framework defines a multi-dimensional design space—the social machine morphospace—within which, it is suggested, all social machines (extant or otherwise) can be accommodated.

The effort to develop a taxonomic framework is important for a number of reasons. Aside from the rather obvious sense in which a taxonomic framework improves our understanding of the similarities and differences between social machines, a taxonomic framework can help us to identify unexplored or under-explored regions of the design space. It can also help to identify clusters of social machines that denote conceptually-important classes or categories of social machines. Finally, the taxonomic framework provides a set of variables that can be exploited in the context of more empirical efforts. For example, some of the dimensions may serve as the independent variables for experimental simulations undertaken as part of cognitive social simulation [47] and computational social science [15] studies.

The research that is reported here forms part of a larger effort to establish a conceptual foundation for social machine research. Given that the recent growth and expansion of the Internet, particularly the Web, has been driven by the emergence of systems such as Facebook, Twitter, YouTube and so on, all of which have been regarded as social machines, the study of social machines is of crucial importance to members of the Web and Internet Science community. In addition, the next generation of social machine systems have been implicated in a range of advanced

capabilities, including curing diseases, solving world hunger, and deriving strategies to mitigate the effects of climate change [18]. This makes the study of social machines of interest to those concerned with our future individual and collective problem-solving capabilities. Finally, social machines are of critical interest in terms of understanding the relationship between the Web and wider society. By supporting the emergence of new forms of social interaction, organization and coordination, social machines are progressively altering the way a broad array of social activities are performed, ranging from the way we communicate and transmit knowledge, establish romantic partnerships, generate ideas, produce goods and maintain friendships. This establishes the basis for more profound forms of social change in which social machines progressively alter the organization and dynamics of our future society. This potential to effect various forms of social change makes the topic of social machines an important focus of research attention for those working across a variety of social science and engineering disciplines.

Acknowledgements This work is supported under SOCIAM: The Theory and Practice of Social Machines. The SOCIAM Project is funded by the UK Engineering and Physical Sciences Research Council (EPSRC) under grant number EP/J017728/1 and comprises the Universities of Southampton, Oxford and Edinburgh.

Appendix

Tables 1, 2, 3, 4, 5, and 6 present the dimensions of the taxonomic framework for social machines.

Table 1 Mapping of social machine dimensions to the dimensions associated with four other systems (i.e., social computing, collective intelligence, human computation and crowdsourcing systems)

Type ^a	Social machine dimension	Target system type	Target dimension	Source
P	Input validation mechanism	Human computation	Quality control	[34]
P	Input validation mechanism	Crowdsourcing	How to evaluate inputs	[11]
P	Human ability	Human computation	Human skill	[34]
P	Human ability	Crowdsourcing	What users can do	[11]
P	Combinatorial strategy	Social computing	From conveyance to convergence content generation	[1]
P	Combinatorial strategy	Human computation	Aggregation	[34]
P	Combinatorial strategy	Crowdsourcing	Aggregation of contributions	[14]
P	Combinatorial strategy	Collective intelligence	How (structure/process)	[25]
P	Combinatorial strategy	Crowdsourcing	How to combine inputs	[11]
P	Task type	Collective intelligence	How (structure/process)	[25]
P	Task type	Crowdsourcing	Type of target problem	[11]
U	Control flow	Human computation	Process order	[34]
U	Task user cardinality	Human computation	Task request cardinality	[34]
U	Sociality	Social computing	From information to people connections	[1]
U	Community specification	Crowdsourcing	Preselection of contributors	[14]
U	Community specification	Collective intelligence	Who (staffing)	[25]
U	Visibility of user contributions	Crowdsourcing	Accessibility of peer contributions	[14]
U	Visibility of user contributions	Crowdsourcing	Nature of collaboration	[11]

U	Response to user contributions	Crowdsourcing	Accessibility of peer contributions	[14]
U	Task assignment policy	Crowdsourcing	Preselection of contributors	[14]
U	Task assignment policy	Collective intelligence	Who (staffing)	[25]
U	Group/individual assignment	Collective intelligence	Who (staffing)	[25]
Q	Quality assessment mechanism	Human computation	Quality control	[34]
Q	Quality assessment mechanism	Crowdsourcing	How to evaluate inputs	[11]
M	Form of motivation	Human computation	Motivation	[34]
M	Form of motivation	Social computing	From utilitarian to hedonic use	[1]
M	Form of motivation	Collective intelligence	Why (incentives)	[25]
M	Reward type	Human computation	Motivation	[34]
M	Reward type	Collective intelligence	Why (incentives)	[25]
M	Reward type	Crowdsourcing	How to recruit and retain users	[11]
M	Reward variability	Crowdsourcing	Remuneration for contributions	[14]

^aIndicates the type of the dimensions from the social machine taxonomic framework: P = 'goal, task and process dimensions', U = 'user participation and interaction dimensions', Q = 'quality assessment dimensions', and M = 'motivational factors and incentive mechanism dimensions'

Table 2 Dimensions and characteristics for the category ‘motivational factors and incentive mechanisms’

Dimension	Description	Characteristics
Motivation type	Specifies the type of motivation associated with user participation	Intrinsic/extrinsic
Form of motivation	Specifies the form of motivation associated with user participation	Economic/altruistic/hedonic/reputational/instrumental/other
Reward type	Specifies the type of reward made for user contributions	None/monetary payment/prize/other
Reward variability	Specifies whether reward quantities are fixed or variable. Variable rewards are encountered when rewards are related to individual or collective performance	Fixed/variable/none

Table 3 Dimensions and characteristics for the category ‘technology and engineering’

Dimension	Description	Characteristics
Open source status	Specifies whether the technological elements of the social machine are open source	Open source/not open source
Social machine framework status	Specifies whether the social machine is derived from a generic framework, such as MediaWiki, Diaspora or Ushahidi	Based on framework/not based on framework

Table 4 Dimensions and characteristics for the category ‘goal, task and process’

Dimension	Description	Characteristics
Goal variability	Indicates whether the goal of the social machine is stable across the lifetime of the social machine, or whether it is likely to change	Fixed/variable
Goal visibility	Indicates whether the goal is visible to the human users of the system	Visible/hidden
Output type	Specifies the kind of output that results from the processes performed by the social machine	Physical/social/cognitive/informational
Output ownership	Indicates who owns the results of process execution	System designer/larger community
Task type	Specifies the kind of task that is performed by the system	Evaluating/organizing/sharing/networking/creating/other
Human ability	Specifies the nature of the primary human ability that is required as part of the process	Aesthetic/emotional/epistemic/perceptual/behavioural/social/moral/cognitive/linguistic
Combinatorial strategy	Specifies how the contributions of individual participants are combined during the course of process execution ^a	Additive/compensatory/disjunctive/conjunctive/discretionary
Input validation mechanism	Indicates how individual user contributions are checked or validated	Automatic/manual/none

^aThese characteristics are derived from Steiner’s [46] categories of task independence

Table 5 Dimensions and characteristics for the category ‘quality assessment’

Dimension	Description	Characteristics
Quality assessment mechanism	Indicates how the quality assessment process is undertaken	Automatic/manual/mixed/none
Explicit/implicit nature of quality criteria	Indicates whether quality assessment criteria are explicitly or implicitly specified	Explicit/implicit
User involvement in quality evaluation	Indicates whether users are involved in the evaluation of process outcomes	User involvement/no user involvement
Quality criteria variability	Indicates whether quality assessment criteria are fixed or variable over the lifetime of the social machine	Fixed/variable

Table 6 Dimensions and characteristics for the category ‘participation and interaction’

Dimension	Description	Characteristics
Social role differentiation	Indicates whether or not users have different roles within the system	Social role differentiation/ no social role differentiation
Functional role variability	Indicates whether or not users are engaged in different processes or the same process as part of their participation in the machine	Functional role variability/ no functional role variability
User autonomy	Indicates the extent to which users decide what they work on and when they work on it	User autonomy/no user autonomy
Community specification	Indicates whether the user community of the system is a subset of the total population. A subset of users may be based on a variety of characteristics, such as demographic factors or the possession of particular skills and abilities	Specified/unspecified
Task atomicity	Indicates whether the user engages in atomic tasks, multiple tasks of the same kind or a combination of tasks	Atomic/multiple instance/combined
Control flow	Indicates the order in which the tasks performed by multiple agencies are executed	Sequence/parallel/split/ synchronization/asynchronous merge/exclusive choice/ iteration
Visibility of user contributions	Specifies the visibility of user contributions to other users of the system	Restricted/unrestricted/variable
User anonymity	Indicates the extent to which participating users are required to provide personal information about themselves to other users	High anonymity/low anonymity
Response to user contributions	Specifies the kinds of ways in which users respond to the contributions made by other users. User contributions may be enriched (e.g., via tagging) or modified. In addition, one user may respond to the contribution of another user by posting related content	None/enrich/modify/respond
User process awareness	Indicates the extent to which users have full knowledge of what is going on in the system	Local awareness/global awareness
Task assignment policy	Specifies how tasks are assigned to users of the system	Random/role-based/skill-based/contribution-based
Task-user cardinality	Specifies the relationship between specific tasks and user assignments	One-to-one/one-to-many/ many-to-many/many-to-one

(continued)

Table 6 (continued)

Dimension	Description	Characteristics
Group/individual assignment	Specifies whether tasks are assigned to individuals or groups	Individual/group
Proportion of active participants	Specifies the proportion of participants that are actively involved in a process as opposed to those who merely consume the contributions of others	High/low/balanced
Sociality	Indicates the extent to which the system supports social interaction with other members	High sociality/low sociality

References

1. Ali-Hassan, H., Nevo, D.: Identifying social computing dimensions: A multidimensional scaling study. In: *Intelligent Conference on Information Systems*. Phoenix, Arizona (2009)
2. Bailey, K.D.: *Typologies and Taxonomies: An Introduction to Classification Techniques*. Sage, Thousand Oaks (1994)
3. Berners-Lee, T., Fischetti, M.: *Weaving the Web: The Original Design and Ultimate Destiny of the World Wide Web*. Harper Collins, New York (1999)
4. Bernstein, A., Klein, M., Malone, T.W.: Programming the global brain. *Commun. ACM* **55**(5), 41–43 (2012)
5. Bonabeau, E.: Decisions 2.0: The power of collective intelligence. *MIT Sloan Manag. Rev.* **50**(2), 45–52 (2009)
6. Caputi, P., Reddy, P.: A comparison of triadic and dyadic methods of personal construct elicitation. *J. Constr. Psychol.* **12**(3), 253–264 (1999)
7. Chi, E.H., Bernstein, M.S.: Leveraging online populations for crowdsourcing. *IEEE Internet Comput.* **16**(5), 10–12 (2012)
8. Clark, A.: *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. Oxford University Press, New York (2008)
9. Clark, A., Chalmers, D.: The extended mind. *Analysis* **58**(1), 7–19 (1998)
10. De Roure, D., Hooper, C., Meredith-Lobay, M., Page, K., Tarte, S., Cruickshank, D., De Roure, C.: Observing social machines Part 1: What to observe? In: *WWW2013 Workshop: The Theory & Practice of Social Machines*. Rio de Janeiro, Brazil (2013)
11. Doan, A., Ramakrishnan, R., Halevy, A.Y.: Crowdsourcing systems on the World Wide Web. *Commun. ACM* **54**(4), 86–96 (2011)
12. Evans, M.B., O’Hara, K., Tiropanis, T., Webber, C.: Crime applications and social machines: Crowdsourcing sensitive data. In: *WWW2013 Workshop: The Theory & Practice of Social Machines*. Rio de Janeiro, Brazil (2013)
13. Fransella, F., Bell, R., Bannister, D.: *A Manual for Repertory Grid Technique*, 2nd edn. Wiley, Chichester (2003)
14. Geiger, D., Seedorf, S., Schulze, T., Nickerson, R.C., Schader, M.: Managing the crowd: Towards a taxonomy of crowdsourcing processes. In: *Americas Conference on Information Systems*. Detroit, Michigan (2011)
15. Gilbert, N., Troitzsch, K.G.: *Simulation for the Social Scientist*, 2nd edn. Open University Press, Maidenhead (2005)

16. Gilles, D., Guattari, F.: *Anti-Oedipus*. Continuum, London (2004)
17. Hall, W., Tiropanis, T.: Web evolution and web science. *Comput. Netw.* **56**, 3859–3865 (2012)
18. Hendler, J., Berners-Lee, T.: From the Semantic Web to social machines: A research challenge for AI on the World Wide Web. *Artif. Intell.* **174**, 156–161 (2010)
19. Jankowicz, D.: *The Easy Guide to Repertory Grids*. Wiley, Chichester (2003)
20. Kelly, G.A.: *The Psychology of Personal Constructs*. W.W. Norton and Company, New York (1955)
21. Klein, G., Moon, B., Hoffman, R.R.: Making sense of sensemaking 2: A macrocognitive model. *Intell. Syst.* **21**(5), 88–92 (2006)
22. Kraut, R., Maher, M.L., Olson, J., Malone, T.W., Pirolli, P., Thomas, J.C.: Scientific foundations: A case for technology-mediated social-participation theory. *Computer* **43**(11), 22–28 (2010)
23. Landes, D.: *Revolution in Time: Clocks and the Making of the Modern World*. Viking Press, London (2000)
24. Law, E., von Ahn, L.: Human computation. *Synth. Lect. Artif. Intell. Mach. Learn.* **5**(3), 1–121 (2011)
25. Malone, T.W., Laubacher, R., Dellarocas, C.: The collective intelligence genome. *MIT Sloan Manag. Rev.* **51**(3), 21–31 (2010)
26. McBride, N.: From social machine to social commodity: Redefining the concept of social machine as a precursor to new Web development approaches. In: 3rd International Conference on Web Science, Koblenz (2011)
27. Meira, S.R.L., Burégio, V.A.A., Nascimento, L.M., Figueiredo, E., Neto, M., Encarnação, B., Garcia, V.C.: The emerging web of social machines. In: 35th Annual Computer Software and Applications Conference (COMPSAC), pp. 26–27. IEEE, Munich (2011)
28. Menary, R.: *The Extended Mind*. MIT Press, Cambridge (2010)
29. Nickerson, R., Muntermann, J., Varshney, U., Isaac, H.: Taxonomy development in information systems: Developing a taxonomy of mobile applications. In: European Conference in Information Systems, Verona (2009)
30. O’Hara, K.: Trust in social machines: The challenges. In: AISB/IACAP World Congress 2012: Social Computing, Social Cognition, Social Networks and Multiagent Systems, Birmingham (2012)
31. O’Hara, K.: Social machine politics are here to stay. *IEEE Internet Comput.* **17**(2), 87–90 (2013)
32. Parameswaran, M., Whinston, A.B.: Research issues in social computing. *J. Assoc. Inf. Syst.* **8**(6), 336–350 (2007)
33. Potthast, M., Stein, B., Gerling, R.: Automatic vandalism detection in Wikipedia. In: 30th European Conference on Information Retrieval Research. Glasgow, Scotland (2008)
34. Quinn, A., Bederson, B.: Human computation: A survey and taxonomy of a growing field. In: Annual Conference on Human Factors in Computing Systems (CHI’11). Vancouver, British Columbia (2011)
35. Rainie, L., Wellman, B.: *Networked: The New Social Operating System*. MIT Press, Cambridge (2012)
36. Raup, D.M.: Geometric analysis of shell coiling: General problems. *J. Paleontol.* **40**, 1178–1190 (1966)
37. Richardson, D., Smart, P.R., Sycara, K., Stone, P., Giammanco, C., Powell, G.: Using ACT-R to model collective sensemaking in military coalition environments. In: Annual Fall Meeting of the International Technology Alliance. Palisades, New York (2013)
38. Robertson, D., Giunchiglia, F.: Programming the social computer. *Philos. Trans. R. Soc. A Math. Phys. Eng. Sci.* **371**(1987), 20120379 (2013)
39. Secretan, J.: Stigmergic dimensions of online creative interaction. *Cogn. Syst. Res.* **21**, 65–74 (2013)
40. Secretan, J., Beato, N., D’Ambrosio, D., Rodriguez, A., Campbell, A., Folsom-Kovarik, J., Stanley, K.: Picbreeder: A case study in collaborative evolutionary exploration of design space. *Evol. Comput.* **19**(3), 373–403 (2011)

41. Shadbolt, N.R.: Knowledge acquisition and the rise of social machines. *Int. J. Hum. Comput. Stud.* **71**(2), 200–205 (2013)
42. Shadbolt, N.R., Smart, P.R.: Knowledge elicitation: Methods, tools and techniques. In: Wilson, J.R., Sharples, S. (eds.) *Evaluation of Human Work*, 4th edn. CRC Press, Boca Raton (2015) (in press)
43. Shadbolt, N., Smith, D.A., Simperl, E., Van Kleek, M., Yang, Y., Hall, W.: Towards a classification framework for social machines. In: *WWW2013 Workshop: The Theory & Practice of Social Machines*, Rio de Janeiro (2013)
44. Smart, P.R.: The web-extended mind. *Metaphilosophy* **43**(4), 426–445 (2012)
45. Smart, P.R., Shadbolt, N.R.: Social machines. In: Khosrow-Pour, M. (ed.) *Encyclopedia of Information Science and Technology*. IGI Global, Hershey, 3rd edn. (2014)
46. Steiner, I.D.: *Group Processes and Productivity*. Academic, New York (1972)
47. Sun, R.: Cognitive social simulation incorporating cognitive architectures. *Intell. Syst.* **22**(5), 33–39 (2007)
48. Surowiecki, J.: *The Wisdom of Crowds: Why the Many are Smarter than the Few*. Random House, New York (2005)
49. Sutton, J.: Exograms and interdisciplinarity: History, the extended mind, and the civilizing process. In: Menary, R. (ed.) *The Extended Mind*. MIT Press, Cambridge (2010)
50. Tinati, R., Carr, L.: Understanding social machines. In: *ASE/IEEE International Conference on Social Computing and International Conference on Privacy, Security, Risk and Trust*, Amsterdam (2012)
51. Tiropanis, T., Hall, W., Shadbolt, N., De Roure, D., Contractor, N., Hendler, J.: The web science observatory. *IEEE Intell. Syst.* **28**(2), 100–104 (2013)
52. Tong, S., Van Der Heide, B., Langwell, L., Walther, J.: Too much of a good thing? The relationship between number of friends and interpersonal impressions on Facebook. *J. Comput. Mediat. Commun.* **13**(3), 531–549 (2008)
53. Tyszk, J.: Morphospace of foraminiferal shells: Results from the moving reference model. *Lethaia* **39**(1), 1–12 (2006)
54. Van Kleek, M., Smith, D.A., Hall, W., Shadbolt, N.R.: “The crowd keeps me in shape”: Social psychology and the present and future of health social machines. In: *WWW2013 Workshop: The Theory & Practice of Social Machines*, Rio de Janeiro, Brazil (2013)
55. von Ahn, L., Dabbish, L.: Designing games with a purpose. *Comm. ACM* **51**(8), 58–67 (2008)
56. Vrandečić, D.: Ontology evaluation. In: Staab, S., Studer, R. (eds.) *Handbook of Ontologies*. International Handbook on Information Systems, 2nd edn., pp. 293–314. Springer, Berlin (2009)
57. Westerman, D., Spence, P., Van Der Heide, B.: A social network as information: The effect of system generated reports of connectedness on credibility on Twitter. *Comput. Hum. Behav.* **28**, 199–206 (2012)
58. Wheeler, M., Clark, A.: Genic representation: Reconciling content and causal complexity. *Br. J. Philos. Sci.* **50**(1), 103–135 (1999)
59. Wilson, R.A., Craver, C.F.: Realization: Metaphysical and scientific perspectives. In: Thagard, P. (ed.) *Philosophy of Psychology and Cognitive Science*. North-Holland, Oxford (2007)
60. Zook, M., Graham, M., Shelton, T., Gorman, S.: Volunteered geographic information and crowdsourcing disaster relief: A case study of the haitian earthquake. *World Med. Health Policy* **2**(2), 7–33 (2010)