

Understanding Human-Machine Networks: A Cross-Disciplinary Survey

MILENA TSVETKOVA, TAHA YASSERI, and ERIC T. MEYER, Oxford Internet Institute, University of Oxford

J. BRIAN PICKERING, VEGARD ENGEN, and PAUL WALLAND, IT Innovation Center, University of Southampton

MARIKA LÜDERS and ASBJØRN FØLSTAD, SINTEF

GEORGE BRAVOS, Athens Technology Center

In the current hyperconnected era, modern Information and Communication Technology (ICT) systems form sophisticated networks where not only do people interact with other people, but also machines take an increasingly visible and participatory role. Such Human-Machine Networks (HMNs) are embedded in the daily lives of people, both for personal and professional use. They can have a significant impact by producing synergy and innovations. The challenge in designing successful HMNs is that they cannot be developed and implemented in the same manner as networks of machines nodes alone, or following a wholly human-centric view of the network. The problem requires an interdisciplinary approach. Here, we review current research of relevance to HMNs across many disciplines. Extending the previous theoretical concepts of socio-technical systems, actor-network theory, cyber-physical-social systems, and social machines, we concentrate on the interactions among humans and between humans and machines. We identify eight types of HMNs: public-resource computing, crowdsourcing, web search engines, crowdsensing, online markets, social media, multiplayer online games and virtual worlds, and mass collaboration. We systematically select literature on each of these types and review it with a focus on implications for designing HMNs. Moreover, we discuss risks associated with HMNs and identify emerging design and development trends.

CCS Concepts: • **General and reference** → **Surveys and overviews**; • **Information systems** → **Web applications**; • **Social and professional topics** → **Management of computing and information systems**; • **Applied computing** → **Law, social and behavioral sciences**;

Additional Key Words and Phrases: Crowdsourcing, mass collaboration, crowdsensing, social media, peer-to-peer, complex networks, human-machine networks

ACM Reference Format:

Milena Tsvetkova, Taha Yasseri, Eric T. Meyer, J. Brian Pickering, Vegard Engen, Paul Walland, Marika Lüders, Asbjørn Følstad, and George Bravos. 2017. Understanding human-machine networks: A cross-disciplinary survey. *ACM Comput. Surv.* 50, 1, Article 12 (April 2017), 35 pages.
DOI: <http://dx.doi.org/10.1145/3039868>

This project has received funding from the European Union's Horizon 2020 research and innovation program under grant agreement No 645043.

Authors' addresses: M. Tsvetkova, Department of Methodology, London School of Economics and Political Science, Houghton Street, London WC2A 2AE, UK; email: m.tsvetkova@lse.ac.uk; T. Yasseri and E. T. Meyer, Oxford Internet Institute, University of Oxford, 1 St Giles, Oxford OX1 3JS, UK; emails: {taha.yasseri, eric.meyer}@oii.ox.ac.uk; J. B. Pickering, V. Engen, and P. Walland, IT Innovation Centre, University of Southampton, Gamma House, Enterprise Road, Southampton SO16 7NS, UK; emails: {jbp, ve, pww}@it-innovation.soton.ac.uk; M. Lüders, Department of Media and Communication, University of Oslo, Postboks 1093 Blindern, 0317 Oslo, Norway; email: marika.luders@media.uio.no; A. Følstad, SINTEF, Forskningsveien 1, 0373 Oslo, Norway; email: Asbjorn.Folstad@sintef.no; G. Bravos, Athens Technology Center, Chalandri 152 33, Greece; email: gebravos@gmail.com.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2017 ACM 0360-0300/2017/04-ART12 \$15.00

DOI: <http://dx.doi.org/10.1145/3039868>

1. INTRODUCTION

Since the invention of the Gutenberg press, machines have had a transformative impact on human interaction. As a result of the Industrial Revolution, machine influence on how we communicate, exchange, and cooperate with each other has become more immediate and pervasive. Today, with the ever-increasing proliferation of networked technology in support of human and machine interaction, human actions and interactions have become so interrelated with technology that it is difficult to know whether society changes because of technology or the other way around.

Increasingly, the interactions of humans and machines form interdependent networks. We conceptualize these networks as Human-Machine Networks (HMNs), that is, assemblages of humans and machines that interact to produce synergistic effects. Ever more activities in work and private life are conducted within HMNs. For example, actions to address environmental problems are executed in networks involving government, interest organizations, citizens, smart devices, and sensor networks. Systems for emergency response and rescue involve complex interactions between sensors, smart machines, victims, volunteers, and emergency response organizations. Scientific research is increasingly conducted in networks of scientists, crowdsourced volunteers and amateurs, and networked machine resources, and many more examples.

In consequence, HMNs increasingly influence our society. For the individual worker and citizen, the form, experience, and outcome of life may sometimes depend less on the characteristics of the individual and more on the characteristics of the online and offline networks in which they are embedded. For companies, the public sector, and organizations, productivity, innovativeness, and civic participation depend on the characteristics of the networks of which workers and citizens are part. Hence, an advanced understanding of HMNs, and how to benefit from their characteristics, is important to strengthen productivity and innovation, as well as citizen participation and quality of life.

The challenge is that HMNs cannot be developed and implemented in the same manner as networks of machine nodes alone. Creating successful solutions for HMNs requires awareness concerning the kind of network to be established, the actors that are involved, the interactions, and the capabilities and behaviors emerging within the network. By establishing such awareness, we may benefit from the existing knowledge and experience when designing the HMNs of the future.

There are four major fields of research of relevance to our understanding of HMNs: Socio-Technical Systems (STS) theory, Actor-Network Theory (ANT), Cyber-Physical-Social Systems (CPSS) theory, and the emerging theory on social machines. STS theory provides insight into the dual shaping of technology and the social (work) context in which it is implemented [Leonardi 2012], recognizing organizations as complex systems of humans and technology that aim to reach given goals in the context of a given organizational environment [Emery and Trist 1965]. ANT aims at fundamentally revising what should be considered part of the social, and argues that we explicitly need to take into account that any social system is an association of heterogeneous elements such as humans, norms, texts, devices, machines, and technology [Latour 2005; Law 1992], granting equal weight to humans and nonhuman (machine) entities in the analysis of the social. Hence, whereas STS theory considers the social as part of technology, ANT considers technology (and other nonhuman entities) as part of the social. More recently, CPSS theory has extended the STS approach and called attention to the human and social dimensions present in systems that use computational algorithms to control or monitor physical devices [Wang 2010; Sheth et al. 2013]. Representing a fourth perspective, the emerging theory on social machines refers to systems that combine social participation with machine-based computation [Burégio et al. 2013; Shadbolt

et al. 2013; Smart et al. 2014], connecting with Berners-Lee's original vision of the Web more as a social creation than a technical one, "to support and improve our Web-like existence in the world" [Berners-Lee 2000, p. 123].

However, even if these theories conceptualize humans and machines as forming a single system, as opposed to perspectives where social structures are seen as merely mediated in machine networks (e.g., social network theory), they do not provide the insight and guidance needed to support the design of contemporary and future HMNs. Furthermore, these theories do not contribute to a unified framework for understanding HMNs, as is seen from the plethora of theoretical positions in which current studies of such networks are based. As will be seen in Section 3, the volume of publications concerning HMNs is rapidly increasing across a wide range of academic fields. Such breadth in academic attention is beneficial for the development of knowledge on HMNs. However, this breadth also implies a potential risk for fragmentation, barring opportunities for the transfer of knowledge and experiences across academic as well as practical fields of study.

To initiate the broad theoretical basis needed to bridge the rapidly evolving field of HMNs, we have conducted a cross-disciplinary literature survey of key studies that concern different types of HMNs. Multiple surveys of particular types of HMNs exist [Crowston et al. 2012; Guo et al. 2013; Yahyavi and Kemme 2013; Guo et al. 2015; Pejovic and Musolesi 2015] but we are the first to present a comprehensive overview of the field within a unifying framework. The objective of our article is to systematically review the existing literature, with a specific focus on the human and machine actors and the interactions between them. This analysis provides a first step toward integrating current academic efforts to understand the emerging phenomenon of HMNs. Furthermore, the analysis demonstrates how the identification of key characteristics in different HMNs may serve to expose issues related to the design of HMNs, as well as to support the transfer of knowledge and experience across academic disciplines and types of HMNs.

In Section 2, we discuss the scope and approach to the systematic literature review we have conducted. In Section 3, we present real-world examples and insights from the literature on different types of HMNs. We discuss common challenges, emergent behavior, and trends in Section 4. We conclude in Section 5.

2. BACKGROUND AND SCOPE

In this section, we provide a definition of an HMN, discuss background literature on the actors that make up HMNs and their interactions, and establish the scope and approach taken to conduct this survey.

2.1. Definition

HMNs are assemblages of humans and machines whose interactions have synergistic effects. This means that the effects generated by the HMN should exceed the combined effort of the individual actors involved. In particular, the HMN should serve a purpose that would not have been achieved merely by the effects of the individual actors. Typically, the synergistic effect is achieved by combining human strengths and machine strengths. Interactions in the HMNs are expected to result in outputs that neither a pure social network nor a computer network could achieve independently.

The concept of HMNs is to be understood as a perspective for analysis or conceptualization. Given our concern for synergistic effects, the networks of main interest are those where the synergistic effects between humans and machines are immediately evident such as in systems for mass collaboration (see Section 3.8). Networks of lesser interest include, for example, simple communication and broadcasting networks such as telephone, telegraph, e-mail, or TV networks, as the medium here may be said to

hold more the role of a noninterfering intermediary. In other words, for a network to be considered an HMN, the machines need to transform and/or influence, not just transmit. In the parlance of actor-network theory, we focus on networks where machines serve as mediators, by affecting and transforming the interaction that takes place.

2.2. Actors

The two types of actors in HMNs are the human and the machine. By “human” we mean an entity that behaves like a single person (even if the entity is an organization). In contrast to machines, humans have the capacity for emotions, attitudes, sociality, meaning-making, creativity, and intent. Compared with machines, human behavior is typically more unpredictable, yet susceptible to influence.

In most cases, any person can participate in an HMN. Participation is more often determined by the benefits and costs the individual perceives than the individual’s social status or geographic location. The benefits can be divided into economic, intrinsic, and social [Ardichvili 2008; Chiu et al. 2006; Yee 2006]. Individuals profit economically when they receive payments (in money, goods, or services) or obtain skills and qualifications for their paid career. They gain intrinsically when they contribute to a good cause, share knowledge, or help others. Some are also intrinsically motivated to gain or create knowledge for its own sake. Individuals may also benefit socially if participation in the HMN allows them to create, build, and maintain reputation and social capital, or simply enjoy social interaction. In general, people often have complex motivations to participate and contribute, and motivation strategies such as competition and gamification make use of this by offering a combination of economic, intrinsic, and social benefits. Time and effort are the two obvious costs that people incur when they participate in HMNs. Another significant cost is the risk of breached privacy [Krasnova et al. 2010; Sheehan 2002]. Leaking private information to the public and stealing a user’s identity can entail significant financial losses, family distress, and social embarrassment.

Regarding machines, any device with connectivity can be part of an HMN: personal computers, smartphones, tablets, wearable technology, sensors, embedded chips, servers running algorithms, and so on. These machines take input in the form of text, media, or sensor data, such as vital signs or environmental measures, printed or coded instructions and software, signals and alerts from multiple sources, including other machines as well as humans. This input can be entered by the user or collected automatically. A machine with computational capabilities can aggregate, clean, or otherwise transform the input data in order to output something else or even fuse data from different sources to make complex decisions to be interpreted by other machines or by human experts [Atzori et al. 2010; Castells 2011; Lee et al. 2015].

To facilitate the functioning of the HMN, the machines need to be available, connected, and secure. Compared to humans, the machines do not have motivation, do not experience trust or reliance, and do not behave altruistically or irrationally of their own volition: they have no agency in that sense [Jia et al. 2012; Rose et al. 2005]. With autonomic computing, however, there may be a promise of some level of self-regulation and “healing” [Huebscher and McCann 2008; Kephart and Chess 2003; White et al. 2004]. Moreover, machines are increasingly capable of solving complex data analysis problems that humans would struggle to tackle not least at such speed, although effective responses and the subtlety of nuanced behaviors remain elusive [Norman et al. 2003; Pantic and Rothkrantz 2003; Picard et al. 2001]. Further, machine learning and data mining crucially depend on large datasets generated and annotated by humans. In such cases, a quasi-balanced collaboration is required. Machines are therefore now becoming increasingly significant contributors to HMNs, and no longer simply communication channels or facilitators of highly distributed human-to-human connectivity.

2.3. Interactions

Machines enable new forms of interaction between humans. Different HMNs allow participants to interact with different levels of intensity and involvement. In some networks (e.g., crowdsourcing platforms such as Amazon Mechanical Turk), participants do not have the opportunity to directly interact. In others, they can observe each other's contributions. Depending on the application, they can browse all contributions (as in an online market such as eBay), see algorithmically selected contributions (as in Reddit, which is a content site), see algorithmically customized contributions (e.g., via the Newsfeed algorithm on Facebook), or only view aggregate results (as in prediction markets). Once being able to observe contributions, participants can then evaluate them by approving them (e.g., "likes" on Facebook), approving or disapproving them ("upvotes" and "downvotes" on Reddit), rating them (Amazon product "stars"), or commenting on them. Finally, participants can sometimes modify others' contributions by editing or deleting them. For example, Wikipedia editors can revert edits and Linux developers can fix code contributed by others. Human-human interactions largely depend on trust [Dwyer et al. 2007; Jones and Leonard 2008]. Phenomena such as trolling, cyber-bullying, and cyber-stalking signify the erosion of trust. Human-human interactions also involve social influence [Bond et al. 2012; Muchnik et al. 2013; Onnela and Reed-Tsochas 2010]. Social influence can lead to large-scale behavioral or emotional contagion. This can be both positive, if it results in the spread of pro-social or health-conscious behavior, and negative, if it leads to dangerous herd behavior or unproductive groupthink.

When the role of the machine in an HMN may be seen as that of a mediator (i.e., its output does not equal the input from the other nodes in the HMN), it is relevant to analyze the human-machine interaction. To some degree, human-machine interactions are gated only by the communication protocol used. Input can be provided via different channels including keyboards, data channels, speech, and movement and gesture, actively or even passively. Interactions in HMNs are therefore increasingly multimodal, with different HMNs allowing individuals to contribute or interact in different ways sometimes dictated only by ease of use [Lin 2003; Lin et al. 2003]. Similarly, the output channel may depend on the service or application (signage, broadcasting, experimentation, web searches, and so forth) or on the preferences of the recipient (a recorded message, a printout, or a visual display).

With respect to what humans input, individuals make contributions to different extents. In some networks, users do not have to do anything to contribute data beyond interacting with the service or even simply turning on the service. In these cases, participants contribute passively. Their social interactions, personal characteristics, and site behavior are automatically collected by the service. These data are then analyzed and often used to recommend actions (as some crowdsensing applications do), social contacts (as in LinkedIn), or products (as in Netflix). In other HMNs, contributors need to make the effort and take the time to contribute actively. They contribute by filling out surveys, composing and editing text, writing code, sending sensor measurements, or uploading videos, photos, and news links. In more extreme cases, participants can also modify the content on the HMN or the rules by which the content is collected, managed, and distributed. Consumers then become "prosumers," and can still switch between the two modes as suits the service or their own needs.

With respect to what machines output, different HMNs may introduce different levels of control over the content. Machines may simply list or show the contributions, as is common in online marketplaces and in direct response to a specific query. Alternatively, they may employ complex algorithms to filter and select contributions, as Web search engines do, for example, personalizing or customizing results to reflect previous user behaviors. Or they may even use bots and human intelligence to

modify contributions. For example, Wikipedia uses vandalism-detection algorithms and forums employ censorship algorithms, while content sites like Digg maintain centralized content management.

2.4. Scope and Approach

To narrow the scope of this survey, we simplify the analysis of machine-machine interactions. In particular, when the underlying or enabling machine networks are perceived as one and the same unit by all the relevant nodes of the HMN, we analyze and represent the machine subnetwork as a single node. This is not to say that machine-machine interactions will not be considered; rather, they will be considered only to the extent that is valuable to the analysis at hand.

Thus, our discussion focuses on the human-human and human-machine interactions in the context of contemporary HMNs. There are clearly differences in the relative roles and status of machines and human actors across the different types of network. For instance, there will be cases where the machines act as a “human proxy” to facilitate what is essentially human-to-human relationships. In other cases, they become semiautonomous agents that act alongside and in collaboration with human agents in the network. We base our analysis on the existence and intensity of these interactions and the extent to which they are essential for the functioning of the HMN.

Based on different possible permutations of interaction types, we identify and differentiate between eight types of HMNs: public resource computing, crowdsourcing, Web search engines, crowdsensing, online markets, social media, multiplayer online games and virtual worlds, and mass collaboration (Figure 1). These types are to be seen as an organizing framework for the survey, not as a complete taxonomy of HMNs. We used keywords associated with these types to systematically select relevant literature, with a focus on recent and high-impact articles. Details of our selection procedure are described in the online appendix.

3. ANALYSIS OF HMNS

The number of publications on HMNs, as measured by the number of articles retrieved by our keywords from the Scopus bibliographical database (www.scopus.com), has increased over the last 20 years (Figure 2). More importantly, the pattern of publications has changed over time. During the late 1990s, relatively few publications used the particular terms in their title, keywords, or abstract. The earliest increase was in the eCommerce area (part of the online markets type), where publications increased from 192 in 1999 to 772 in 2000, and remained near or above 1,000 publications per year after that. The second area to gain prominence is file sharing, which increased from less than 50 publications annually in 2000 to approximately 2,000 per year by 2005.

The pattern for social media is particularly striking in the chart, showing a steep rise starting in 2008 ($n=708$) and accelerating by 2010 ($n=2,560$) to a high of 8,991 publications by 2014. Mass collaboration as a topic, on the other hand, shows a more steady rise over the last 10 years, peaking at 5,165 publications in 2014.

We next look at the extent to which ideas relating to human-machine interaction have penetrated various disciplines as indicated by the journals where these publications appear. In Figure 3, we use the methods described in Leydesdorff et al. [2013] and Leydesdorff et al. [2015] to create an overlay map of the journals in the data extracted from Scopus. The method involves extracting the journal names and processing them with Leydesdorff’s tools.¹ These tools match our Scopus data with a standardized map of science generated using journal-journal citations to discover the closeness of journals to each other, and thus the relative location of disciplines. The underlying map thus

¹The tools are available at http://www.leydesdorff.net/scopus_ovl/.

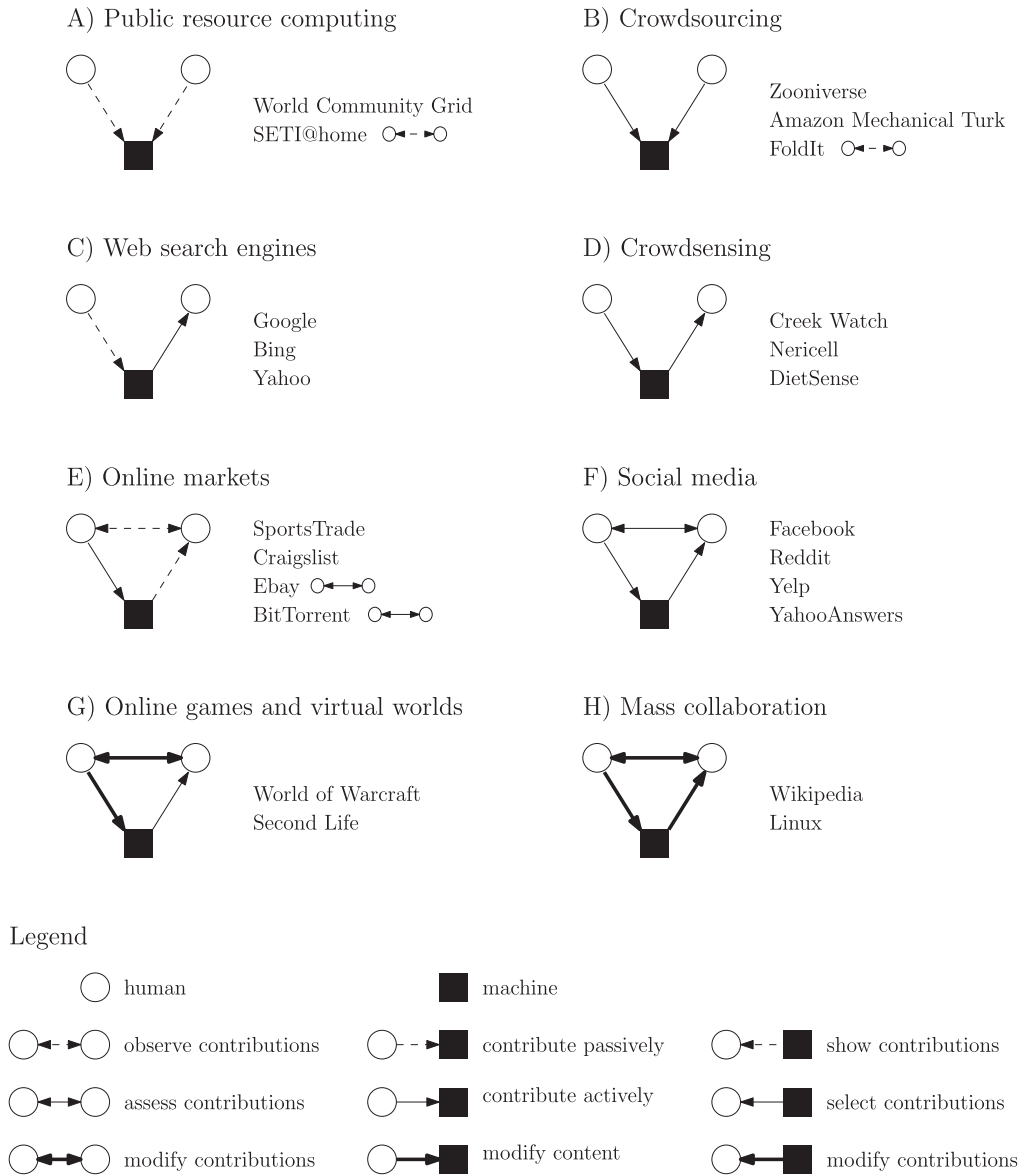


Fig. 1. Types of HMNs.

represents “all of science,” or at least as much of it as is represented in the 19,600 journals indexed by Scopus.

We observe that HMNs, as defined by our keywords, spread across all areas of science. There is a concentration one might expect in the upper left portion of the figure in the computer science-related journals, but there is nearly as much publication activity in the physical sciences at the upper right, the medical sciences at the lower right, and in the social sciences and (to a lesser extent) the humanities in the lower left.

In what follows, we present the most prominent articles from this interdisciplinary literature, related to human participation and to human-human and human-machine interactions.




Fig. 2. Publications related to HMNs.

3.1. Public Resource Computing

Public Resource Computing (PRC) projects are distributed computing networks, where each processing unit is a personal computer voluntarily embedded in the network to collaboratively solve computational problems. The majority of PRC projects are dedicated to scientific projects in fields as diverse as astrophysics, mathematics, molecular biology, and seismology, where extremely computation-intensive tasks are common. PRC networks are the simplest HMNs as far as the human-human and human-machine interactions are concerned (Figure 1(A)). Once volunteers sign up for a PRC project, they are no longer required to actively contribute time or effort apart from what their CPU is contributing. At the same time, they will not receive any substantial feedback from the project either. The first PRC systems were launched in 1996–1997: GIMPS (The Great Internet Mersenne Prime Search) and Distributed.net [Anderson et al. 2002]. SETI@home is another early instance of a PRC network. It was launched in 1999 with the goal to analyze radio signals from space in search of extraterrestrial intelligent life. Further examples include the World Community Grid, intended to tackle scientific research projects in any field [Hachmann et al. 2011] and GPUGRID, designated to molecular dynamics simulation methods [Buch et al. 2010].

The design of PRC systems crucially depends on building the machine-machine interaction infrastructure. Nevertheless, considerations regarding the recruitment and motivation of contributors remain important. To attract participants, successful PRC platforms have emphasized the social value of their scientific projects and let mass media news coverage and word-of-mouth do the rest [Anderson et al. 2002]. To retain participants and increase contributions, they have built reputation systems based on numeric measures of contributed computation and allowed the ability to view others' profiles, participate in message boards, and form and join teams [Anderson 2004; Beberg et al. 2009]. Such an infrastructure stimulates individual and group competition. Empirical research suggests that team competition is particularly important. Being affiliated with a team significantly increases one's contribution. More importantly, being part of a team maintains high levels of contribution over time [Nov et al. 2010].




Fig. 3. Journals publishing articles related to HMNs. The colored dots are Scopus journals that include articles with the key terms we identified; the background uncolored dots do not include such articles. The colors represent broad ‘community’ groupings distinguished by the visualization software. Larger dots with corresponding larger labels have more published articles, and smaller dots/labels have fewer.

As PRC systems have grown in popularity and sophistication, new design challenges have emerged. These have been reflected in the structure of the machine-machine subnetwork. One of the major issues in PRC is managing the expected high rates of machine node unavailability. In current PRC implementations, this has led to redundant computing [Anderson 2004; Mayer et al. 2015]. Redundancy also helps resolve erroneous computational results due to malevolent users or, more often, malfunctioning machines. Alternative solutions to node unavailability include designing a framework core architecture that contains an extra layer that monitors, manages, and brokers resources [Cuomo et al. 2012] and designing a hybrid resource architecture by supplementing volunteer computers with a small set of dedicated, reliable computers [Lin et al. 2010].

Another major problem in current PRC systems is the back-end bottleneck due to the centralized storage of results. This problem can be resolved by combining PRC with public resource storage, thus limiting the volume of data transfer as well [Beberg et al. 2009]. Additional issues of unique authentication, authorization, resource access, and resource discovery can be resolved by borrowing ready-made solutions from grid computing [Foster et al. 2001].

Finally, current implementations of PRC are limited to tasks with independent parallelism, meaning that participating machines are not required to communicate with each other [Anderson et al. 2002]. Newer projects, however, take the form of peer-to-peer networks. For example, Marozzo et al. [2012] develop a MapReduce-based peer-to-peer PRC system in which autonomous nodes are dynamically assigned slave roles, master roles, or backup master roles. In another example, Mayer et al. [2015] implement a Cloud system of a loose set of voluntarily provided heterogeneous machine nodes that uses a gossip-style protocol for communication.

3.2. Crowdsourcing

Jeff Howe originally coined the term crowdsourcing in the June 2006 issue of *Wired* magazine to denote situations where a company or institution take a function once performed by employees, and outsource it to an undefined network of people in the form of an open call [Howe 2006]. The term “crowdsourcing” has since been used in a very broad sense to denote a range of HMNs [Doan et al. 2011; Geiger et al. 2011]. Here, we use it to denote systems that are based on open calls for the voluntary undertaking of tasks [Estelles-Arolas and González-Ladrón-de Guevara 2012], hence closely aligned with Howe’s original definition. Crowdsourcing (CSO) is, as such, a top-down initiated process, and usually distinctly different from mass collaboration as in, for example, open software development (see Section 3.8). The defining characteristics of these HMNs are that humans actively select tasks and contribute to them but do not generally interact and collaborate with each other directly. More importantly, humans rarely receive any feedback from the machine—they do not obtain any information on or direct benefits from the project they contribute to (Figure 1(B)). In a sense, CSO platforms are similar to distributed computing systems: each user is equivalent to a processor that needs to solve a task requiring human intelligence [Kittur et al. 2013].

Prominent examples of CSO HMNs include online microwork platforms such as Amazon Mechanical Turk and Crowdfunder, voluntary mapping websites such as Ushahidi and OpenStreetMap, and citizen science projects such as Zooniverse and FoldIt. Companies and organizations may also crowdsource innovation tasks to the general public (as with Threadless) or experts beyond the organization’s boundaries (such as Innocentive) [Brabham 2008], a practice which is often utilized within the field of open innovation. Perhaps the CSO project that is most familiar to the general public is reCAPTCHA. Web users contribute to reCAPTCHA when they take a Turing test on a website in order to confirm that they are human and not a computer. As part of the test, they transcribe scanned words that optical character recognition programs failed to recognize and thus help digitize large volumes of old text.

The successful functioning of CSO HMNs requires recruiting and maintaining a pool of intelligent, diverse, and capable contributors [Saxton et al. 2013] and managing and processing their contributions. Various approaches can be undertaken to successfully recruit and keep contributors. One way is to make the contributions implicit, as in reCAPTCHA. Most individuals complete the task unbeknownst to them—they are simply trying to register for an online service or make a purchase, for example. Importantly, reCAPTCHA is neither more difficult nor more time-consuming than a regular CAPTCHA test [von Ahn et al. 2008].

Another method to motivate contributors is to remunerate them. Payments can be fixed or success-based. They can come in the form of a piece rate, a piece rate with a bonus, a quota rate, or a contest prize. Experimental research has shown that higher piece rates increase the quantity but not the quality of contributions; further, quota rates result in higher quality work for a smaller budget than piece rates [Mason and

Watts 2010]. Prize-based contests are optimal for highly uncertain innovation problems as they involve a high number of entrants who execute multiple independent trials, which increases the probability for a maximally performing solution [Lakhani et al. 2013]. The effect of payment is not always straightforward, however, as it may reduce social motivation [Frey and Jegen 1999].

Much participation in CSO projects is voluntary and indeed driven by intrinsic and social motivation. To attract more volunteers, CSO platforms can emphasize the connection between people's contributions and the project outcomes. For example, citizen science projects can keep participants updated on the articles published as a result of their contributions [Cooper et al. 2010]. Volunteer-mapping sites can successfully rally on social media around natural disasters [Zook et al. 2010].

Gamification is another strategy to recruit volunteers and encourage their prolonged engagement in a CSO project. A game offers the players a varied motivation set that leverages ambition, competition, and cooperation. Thus, a gamified CSO project should allow for short-term and long-term rewards through tracking game score and player rank, social praise by allowing comments on chats and forums, and collaboration by enabling the formation of teams [Cooper et al. 2010]. These strategies have been successfully implemented in biology projects for discovering protein structures [Cooper et al. 2010], aligning multiple sequences of DNA [Kawrykow et al. 2012], and solving *in vitro* RNA design problems [Lee et al. 2014].

Managing the contributors' work poses the second major challenge in designing CSO platforms. Contributions need to be collected, processed, and aggregated. Although the quantity of contributions is what makes CSO work [Surowiecki 2005], the quality of contributions remains of paramount importance.

One major challenge with CSO is dealing with biases and errors in contributors' submitted work, whether malevolent or accidental. Quinn and Bederson [2011] propose multiple techniques to do this: output agreement, input agreement, economic incentives, defensive task design, reputation systems, redundancy (combined with majority consensus), ground truth seeding, statistical filtering, multilevel review, expert review, and automatic check. Allahbakhsh et al. [2013] add to the list runtime support, workflow management, worker selection, and contributor evaluation. The latter two have been implemented in sophisticated algorithms that assign weights to workers based on the quality of their contributions in order to account for biased responses [Ipeirotis et al. 2010; Raykar et al. 2010] and eliminate responses by spammers [Raykar and Yu 2011].

Beyond correcting for intentional or unintentional errors that have already occurred, one could implement more foresighted strategies to improve the quality of contributions. These strategies usually involve creating and structuring interactions between contributors. For example, Kittur et al. [2011] and Bernstein et al. [2010] implemented frameworks that manage the coordination between users to let them complete complex tasks such as writing and editing articles. Silvertown et al. [2015] purposefully connected participants with experts in a hybrid learning-CSO platform for the identification of biological species. Khatib et al. [2011] introduced tools for free collaboration between contributors that lead to the discovery of new algorithms for the folding of proteins. Bayus [2013] suggests that encouraging contributors to comment on diverse sets of ideas proposed by others could prevent creativity stagnation due to the "fixation effect" and help maintain an ongoing supply of good new ideas in idea-generating CSO platforms.

3.3. Web Search Engines

A Web Search Engine (WSE) is a system that searches and returns content from the Web. Apart from formulating queries, users do not actively contribute to the system. Their contribution is passive: the WSE observes and analyzes their search behavior to

improve and customize search results. Additionally, the search results are filtered and ranked, which is the most important function of these HMNs (Figure 1(C)).

The greatest challenge that WSEs need to address is keeping up with the scale of content on the Web [Arasu et al. 2001; Brin and Page 2012]. Since this problem entails complex technical solutions at the level of machine-machine interactions, it is out of the scope of the present survey. Instead, we here focus on design solutions related to the motivation of users to participate in the improvement of the output of WSEs. These design solutions are typically aimed at making user participation an implicit part of their search task as well as their linking behavior, which is reminiscent of the user motivation strategies of crowdsourcing HMNs such as reCAPTCHA discussed previously.

Another key challenge of WSEs is providing users with relevant search results, quickly. Query assistance is one mechanism introduced to help with this, which is addressing the direct interaction between the users and the WSE interface. One approach is to do this based on the user's own input of a query and subsequent query reformulations [Huang and Efthimiadis 2009]. Alternatively, other users' inputs can be used by mining query associations from multiple users [Carpineto and Romano 2012], which utilizes the HMN further.

To improve the search results, filtering and ranking can also use information about users' behavior. For example, the links between Web pages imply recommendations from users (owners of the respective Web pages) [Arasu et al. 2001]. This implicit notion is encapsulated in Google's PageRank algorithm, for example, as it determines the importance of a page based on the importance of the pages that link to it [Page et al. 1999]. Another well-known ranking algorithm, HITS, also uses recursive logic to differentiate two types of important Web pages, hubs and authorities: hubs are pages that point to many authorities and authorities are pages that are pointed to by many hubs [Kleinberg 1999]. Newer ranking algorithms tend to use even more information about the linking structure. For example, ClusterRank takes into account not only a page's number of linked neighbors and the neighbors' influence scores but also the neighbors' links with each other [Chen et al. 2013].

Gollapudi and Sharma [2009] suggest that WSEs need to intentionally diversify results to improve user satisfaction; simply providing users with the highest ranked pages is not considered sufficient. To personalize search results and improve their relevance, WSEs can also analyze individual behavior [Steichen et al. 2012]. For example, Collins-Thompson et al. [2011] propose algorithms to rerank search results based on an estimation of the user's reading proficiency. Jansen et al. [2008] suggest that WSEs can use knowledge of user intent (whether informational, navigational, or transactional) to provide more relevant results.

Finally, WSEs can also use peers' behavior to improve the results from individual searches. This relatively recent model is known as social search. First, WSEs can potentially incorporate search data from social bookmarking [Heymann et al. 2008]. Second, WSEs can mine the profiles of a user's online social network friends and return results based on their recommendations [Morris et al. 2010]. WSEs can also be adapted to support collaborative search in real time [Morris and Horvitz 2007; Morris 2013].

3.4. Crowdsensing

In crowdsensing (CSE) applications, users with computing and sensing devices (e.g., smartphones or fitness trackers) share data and the application uses the data to measure and map certain phenomena [Ganti et al. 2011]. The contributions by users can be active, as well as passive. The contributions are then analyzed and action and behavior recommendations are sent back to the user (Figure 1(D)).

CSE has been employed in many different areas, including traffic, health, and environmental monitoring [Khan et al. 2013]. Applications have been developed to create

real-time awareness of earthquakes [Faulkner et al. 2014], track and encourage physical activity and healthy lifestyle [Consolvo et al. 2008], track personal transportation patterns to encourage “green” transportation behavior [Froehlich et al. 2009], and predict bus arrival times [Zhou et al. 2012], among many others.

Participating in CSE incurs real costs: increased consumption of energy and bandwidth, as well as increased risk of leaked personal information [Ra et al. 2012]. As a result, most research on CSE has focused on designing monetary incentives for participation, reducing participation load, improving energy efficiency, and guaranteeing privacy.

To guarantee an adequate number of participants, some researchers have proposed monetary incentives through reverse auctions, whereby users claim bid prices for their sensing data [Koutsopoulos 2013; Lee and Hoh 2010]. Others have suggested that this could be achieved through opportunistic sensing, that is, sensing that is fully automated and does not require the user’s active involvement [Lane et al. 2010]. Full automation, however, is difficult to achieve technically, as the application needs to combine data from multiple sensors to infer the context. Instead, others have focused on improving the energy efficiency of the CSE applications. Lu et al. [2010] present methods to turn sensing on and off depending on the quality of input data and the user’s long-term behavior and mobility patterns. Sheng et al. [2012] propose a collaborative sensing approach, whereby data collected from mobile phones is analyzed in real time on servers in a Cloud in order to calculate the best sensing schedule and inform the phones when and where to sense.

Guaranteeing privacy is a more challenging problem to solve. CSE applications are in danger of leaking personal data such as users’ location, speech, potentially sensitive images, or biometric data [Christin et al. 2011]. Even if sensing data look safe, they may be reverse-engineered to reveal invasive information. Last but not least, there is also the “second hand smoke” problem: a person with a sensor can undermine the privacy of nearby third parties [Lane et al. 2010].

Effective privacy-protection measures can be implemented at every stage of the path of the data, from collection to consumption [Christin et al. 2011]. When reporting data, spatial cloaking techniques can be combined with data perturbation to improve user anonymity without affecting aggregate data estimates [Huang et al. 2010]. During that stage, phones can also combine their data with data from their neighbors before transmitting it to the application server [Li et al. 2014]. Techniques like this can be effectively combined to allow even for anonymity-preserving reputation systems [Wang et al. 2013]. However, it has been shown that even if the collected temporally and geographically tagged data are anonymized, only few data points can be uniquely linked to individuals in large crowds [de Montjoye et al. 2013].

Processing and analyzing the collected data presents the second substantial design challenge for CSE systems. Due to user mobility, density, and privacy preferences, the data delivered often have many measurement gaps in both time and space [Ganti et al. 2010]. To be able to generalize from a sparse sampling of high-dimensional spaces, CSE systems can employ sophisticated data interpolation techniques [Mendez et al. 2013]. Further, to make best use of the data a user reports, they can infer the context with the help of machine learning techniques [Pejovic and Musolesi 2015]. Finally, they can combine context-rich data from multiple users to predict models of future behavior and offer action recommendations [Pejovic and Musolesi 2015].

3.5. Online Markets

Under the type “online markets” we group a number of HMNs that involve the exchange of goods and services. The things that are exchanged differ but what unites these HMNs is the common relational structure. In particular, users can view each

other's contributions and, in some cases, evaluate them, but cannot modify them. More importantly, while users actively make contributions, the service does not filter or rank them but simply lists them, either individually or in an aggregate form (Figure 1(E)). In other words, the service usually does not bias the content a user is exposed to.

Online markets are very familiar to most Web users. First, there are consumer-to-consumer markets such as Craigslist, eBay, Uber, and Airbnb, which connect users who offer material goods and services with users who seek those goods and services. File sharing networks such as Gnutella and BitTorrent are similar with the exception that the exchanged goods are digital. Finally, prediction markets, such as the Iowa Electronic Markets and TradeSports, where users trade "stocks" tied to future events, are another prominent example of online markets.

Peer-to-peer (P2P) and file sharing markets involve unregulated direct interuser transactions and as a result, their functioning crucially depends on trust. The problem of trust online is most commonly addressed by designing and implementing reputation systems [Jøsang et al. 2007]. In a reputation system, users rate each other after a transaction and the aggregated ratings inform other users in their own transaction decisions. When designing reputation systems, it is important to take certain well-established empirical facts about social influence into consideration. Previous research has shown that when users deal with unknown sellers or providers, positive reputation matters but negative reputation is much more decisive [Standifird 2001]. The opposite appears to be true when users interact with well-known brands: positive information has a stronger effect on purchasing decisions than negative information [Adjei et al. 2010]. High reputation means high trust by others but who those others are also matters [Baek et al. 2012]. For example, recognition by a well-known other, such as an institution, affects one's trust positively to a higher extent than recognition by one's peers [Jones and Leonard 2008]. The idea to weigh reputation by the trustee's importance in the trusting network has been implemented in two recent reputation algorithms. The EigenTrust algorithm [Kamvar et al. 2003] is based on the rule of transitive trust: if I trust a user I am also likely to trust the people that that user trusts. The PowerTrust algorithm [Zhou and Hwang 2007] relies on the idea of identifying the few power users who can indisputably serve as authorities.

Individuals' estimated trust and reliability could also inform structural redesigns of the online market. For example, Saroiu et al. [2001] propose to use this information to delegate different responsibilities to different users in P2P file sharing applications. Such intervention is intended to mitigate free riding, where most users consume without paying back to the community in return [Hughes et al. 2005].

Information on existing social structure can be used to improve not only trust and cooperation in online markets but also content search. Sripanidkulchai et al. [2003] proposes to preserve the anonymity of the system but group users by interests (not externally visible). Since users who already share interests are likely to have more interests in common, a content location algorithm that exploits the interest overlay network results in faster content queries, lower system load, and improved scalability. Such implicit grouping by similarity in interests and behavior has already been implemented in "collaborative filtering" recommender systems, which are common in online markets [Schafer et al. 1999, 2001]. In contrast, Pouwelse et al. [2008] propose deanonymizing the P2P system by introducing social-network capabilities in order to improve content discovery and recommendation.

Prediction markets differ from P2P and file sharing markets because the exchanges in them are centrally regulated. Consequently, they are less affected by the problem of trust. Instead, the major problem they face is how to improve predictions. Once again, the solution has to do with careful design of the underlying user-to-user interaction structure. Prediction markets work because of the diversity of users [Surowiecki 2005].

This makes them robust to manipulation by a small group of individuals [Hanson et al. 2006; Wolfers and Zitzewitz 2004]. However, social influence can undermine diversity: it both shifts the average estimate and increases the users' confidence in it [Lorenz et al. 2011]. Hence, prediction markets should reduce possibilities for interactions between users, such as observing users' current prediction, past performance, or formal expertise. Nevertheless, to improve the prediction market performance, it is possible to combine the predictions by the crowd with predictions by a panel of experts [Prokesch et al. 2015].

3.6. Social Media

In one of the most cited definitions, Social Media (SM) is defined as “a group of Internet-based applications that build on the ideological and technological foundations of Web 2.0, and that allow the creation and exchange of User Generated Content” [Kaplan and Haenlein 2010, p. 61]. Kaplan and Heinlein's definition is very broad, with the consequence that they include virtual worlds such as Second Life and mass collaboration projects such as Wikipedia as examples of SM. We use a more subtle definition and understanding, focusing on the typical interactions among the nodes in the network. In SM, human nodes assess each other's contributions (rather than modify them) (see Figure 1(F)). SM is intended to enable users to form online communities and interact and share information in them [Kim et al. 2010]. In SM, users can contribute content actively. They are also actively involved in observing and evaluating each other's contributions. The applications use these evaluations to filter the content that users receive. This filtering can be site-wide or user-specific.

SM includes social network applications such as Facebook and LinkedIn. In these applications, users maintain a profile and a list of users with whom they are connected, both of which can be viewed by others [Boyd and Ellison 2007; Donath and Boyd 2004]. SM also includes content communities such as Reddit and Tumblr, where users share news and media and vote and comment on each other's contributions. Some SM applications have the properties of both social network and news-sharing applications, for example, Twitter [Kwak et al. 2010]. Additional examples include review and rating applications such as Yelp and TripAdvisor and discussion forums and question-and-answer sites such as Yahoo Answers and College Confidential.

Participation in SM highly depends on users' motivation and concern for privacy. Based on data from Facebook users, Lin and Lu [2011] found that enjoyment is the strongest motivator to use SM. Still, utilitarian considerations play a role: users often use SM to organize social events and disperse news in an effective way [Xu et al. 2012]. Further, SM is self-affirming, in the sense of satisfying the users' need for self-worth by allowing them to exhibit a successful, attractive, and well-connected version of themselves [Toma and Hancock 2013]. More interestingly, there is a network externality effect on motivation: high levels of adoption of SM among one's peers not only increases the perceived enjoyment of the SM [Lin and Lu 2011] but also directly increases one's likelihood of adoption [Sledgianowski and Kulviwat 2009].

Protecting personal data and privacy is a complex problem in SM. Even if users keep their profiles private, their friendships and group affiliations sometimes remain visible. In addition, some of their friends may have public profiles. Previous research has shown that friendships, group memberships, and rating behavior can be used to infer sensitive personal attributes and information [Kosinski et al. 2013; Zheleva and Getoor 2009]. Users' privacy is threatened by other untrustworthy users, ill-intentioned third parties, as well as the SM providers themselves. To protect users from other users, SM developers can analyze the SM network structure to infer users' trust [Mislove et al. 2007]. Further, they can develop machine learning models to describe a user's privacy preferences toward each of their network contacts; these models can then be

used to configure that user's privacy settings automatically [Fang and LeFevre 2010; Gilbert and Karahalios 2009]. For even more strict privacy, the architecture of the SM application can be decentralized so that users store their private data on other users' machines, such as individuals whom they trust in real life [Cuttillo et al. 2009].

One of the most significant functions of SM is user-specific filtering and customization of content due to the sheer volume of data. While this addresses a practical issue from the users' perspective, it inevitably introduces bias (intended or not) that influences the users. SM employs complex algorithms to predict the news that users would like to read, the media they would like to see, the people they would like to befriend, and the ads they are likely to succumb to. For example, SM applications can analyze user relationships to select high-quality content [Agichtein et al. 2008]. They can also incorporate various behavioral information, such as locational co-occurrences, to recommend new links [Scellato et al. 2011]. SM applications can take even more proactive roles. They can implement algorithms using the political valence of contributions to influence a user's opinion [Bakshy et al. 2015]. They can filter content based on their emotional valence to sway a user's affect [Kramer et al. 2014]. They can also encourage links that cut across traditional age and gender homophily in order to expose users to more diverse sources of information and influence [Centola and van de Rijt 2015]. These strategies can help alleviate one of the major downsides of SM—the potential for knowledge bubbles and misinformation epidemics due to selective interaction.

3.7. Multiplayer Online Games and Virtual Worlds

Multiplayer Online Games (MOGs) and Virtual Worlds (VWs) are simulated environments in which users participate and interact with each other and the environment via avatars that represent their identity. Users' actions and contributions impact other users' experiences both directly and indirectly, by affecting the game world as a whole. In fact, in VWs, it is users who create the game world [Ondrejka 2004]. Apart from providing the interaction rules and settings, the MOG/VW service does not usually modify user contributions (Figure 1(G)).

The question as to why users repeatedly participate in MOGs and VWs has attracted much research attention. As with SM, survey studies have shown that providing optimal personal and social interactions improves users' experience and increases their loyalty [Choi and Kim 2004]. For MOGs, gratification from achievement and social interaction together with game incentives and fairness have been shown to make users more likely to continue playing [Wu et al. 2010]. For VWs, perceived usefulness has been found to additionally increase the likelihood to participate [Verhagen et al. 2012]. Moreover, the functional, experiential, and social motivations appear to be roughly equally important [Zhou et al. 2011]. What is particularly unique for MOGs and VWs, however, is the sense of immersion or "flow" [Yee 2006]. Flow experience is the mental state of being fully absorbed and losing track of time [Goel et al. 2013]. Providing optimal personal experiences and meaningful social interactions cause individuals to experience flow, which in turn increases their likelihood to continue interacting and playing [Choi and Kim 2004; Goel et al. 2013]. Customization (e.g., of one's avatar) further increases loyalty [Teng 2010].

How does one design for meaningful social interactions in MOGs and VWs? Users may be involved in both negative interactions, such as attacks, and positive interactions, such as communication and exchange [Szell et al. 2010]. Overall, however, observational research has shown that MOG users do not interact directly with other users as much as expected [Ducheneaut et al. 2007]. This may be considered suboptimal because, in theory, MOGs and VWs can serve as a venue for informal sociability that fosters bridging social relations, which are relations that improve access to diverse information [Steinkuehler and Williams 2006]. Ducheneaut and Moore [2004] propose

to force social interaction among users through the game design. For example, MOG developers can design interdependencies among characters and then design locations where these interdependencies can play out. Another solution to socially engage users is to design for indirect interactions. Ducheneaut et al. [2006] argue that MOG users do not interact directly and collaborate with other users as much as use them as an audience, a source of entertainment, and a source of information and chatter.

Designing a MOG/VW to keep users playing is a different problem from designing it to keep attracting new users. The more users there are, the more interactive and complex the virtual environment is. In order to guarantee continuous growth of the MOG/VW, one can reduce the learning and personalization costs in the short term [Zhang et al. 2014]. In the long term, however, the problem of scalability looms large. A successful MOG/VW should provide consistent and secure service with fast response times for thousands of users simultaneously [Claypool and Claypool 2006; Yahyavi and Kemme 2013]. Traditional client-server architectures, however, have inherent scalability limitations. In contrast, P2P architectures can achieve high scalability at low infrastructure cost [Hu et al. 2006; Yahyavi and Kemme 2013]. In the context of MOGs and VWs, a P2P architecture means that each user's machine may hold master copies of some of the game objects and be responsible for propagating updates to other nodes. While improving scalability, such an architecture is vulnerable to cheating, security attacks, and churn, problems that remain outstanding design challenges today [Yahyavi and Kemme 2013].

3.8. Mass Collaboration

Compared to the other HMNs, mass collaboration networks involve highly intense interactions among humans and machines in the context of a “project” set up for a specific purpose (Figure 1(H)). They require the highest level of involvement from users in terms of time and effort. Users can modify and reject each other's contributions and affect the project as a whole. Their contributions and the project content are often monitored and adjusted by the project leaders. This centralized oversight can happen automatically, for example, through vandalism-detection algorithms. Nevertheless, often, the organization of these HMNs is bottom-up.

Wikis and Open-Source Software (OSS) projects present two of the most prominent examples of mass collaboration HMNs. Wikis enable the simultaneous collaborative creation, modification, and deletion of content [Tapscott and Williams 2011]. Wikipedia, the online encyclopedia is the largest, most successful, and most popular wiki project. OSS projects allow the collaborative development and free distribution of computer software. Examples include the Firefox web browser, the Apache HTTP Server, and the Linux operating system.

Mass collaboration projects are developed by geographically and organizationally dispersed contributors, most of whom are volunteers. As a result, recruiting and retaining contributors is a critical design issue for OSS and wiki communities. Von Krogh et al. [2003] suggest that “joining scripts” are important if users want to gain access to the collaboration community. These “scripts” are implicit constructs that determine the typical level and type of activity a joiner needs to go through before becoming a contributor. Users often decide to join and participate in a collaboration project because they have pragmatic considerations or expect external rewards. Survey results have shown that external rewards have greater weight for participation than internal factors, such as intrinsic motivation, altruism, and community identification [Hars and Ou 2001]. In fact, a high number of developers are paid for their OSS development efforts; others use their participation to improve their own software; still others receive benefits in terms of reputation and self-development. Contributors are able to learn even from mundane tasks such as reading and answering users' questions [Lakhani and von Hippel 2003].

In addition, contributors can benefit by exchanging valuable work with each other. A theoretical model has shown that since more modular codebases with more option value foster such exchanges, they increase recruitment and retention and decrease free riding [Baldwin and Clark 2006].

Still, external rewards are not the sole motivator. First, they appear to be less important to wiki contributors, compared to OSS contributors [Oreg and Nov 2008]. Second, they appear to be only driving the decision to join but not the decision to stay [Shah 2006]. Sustained participation appears to be better predicted by group identity and community belonging. For example, Hertel et al. [2003] show that self-identification as a Linux developer is one of the factors that determine engagement in the Linux project. Fang and Neufeld [2009] similarly find that situated learning and identity construction are positively linked to sustained participation, at least in the phpMyAdmin OSS community. Bagozzi and Dholakia [2006] confirm that active participation is associated with group-referent intentional actions.

In addition to recruiting and retaining contributors, mass collaboration platforms need to organize the leadership, coordination, and collaboration among the contributors [Crowston et al. 2012]. Analyses of the e-mail communication network of contributors to several OSS projects reveal that subcommunities emerge naturally to mirror collaboration relations [Bird et al. 2008]. The community further subdivides into a core of usually 10–15 developers who create about 80% of the code functionality, a much larger group around the core who repair defects in the code, and an even larger periphery of users who report problems [Mockus et al. 2002]. Members of the core usually know and trust each other and communicate intensely to manage the dependencies among the contributed code.

How do these naturally emerging structures affect the growth and the success of the OSS/wiki project? Crowston and Howison [2005] analyzed 230 project teams on SourceForge to find that larger projects tend to have more decentralized communication networks. Based on a longitudinal analysis of an-order-of-magnitude larger SourceForge sample, Singh et al. [2011] conclude that internal cohesion, as defined by repeat ties, third-party ties, and structural equivalence among contributors, is also associated with project success. Similarly, Hahn et al. [2008] find that a project is likely to attract more developers if prior collaborative relations in the OSS developer network exist. However, contributors' external embeddedness has more complex effects on success. While a high number of external contacts increases success, only moderate technological diversity of the external network and moderate external cohesion are beneficial [Singh et al. 2011]. Further, contributors' participation in multiple projects can both improve and aggravate the project's chances for technical success [Grewal et al. 2006].

4. DISCUSSION

Based on the HMNs reviewed earlier, here we provide further discussion on the emerging trends, common challenges, design implications, and future considerations.

4.1. Emerging Trends

Our survey suggests five major trends that concern most of the HMNs.

First, human-human interactions appear to be intensifying in HMNs. Many systems are starting to allow for social interactions among their users. In addition, many of them are looking for ways to make the social interactions more immediate and intense. In a sense, everything can now become a social network; everything can be “shared.” Social events and meet-ups have been shown to be effective in recruiting new users as well as stabilizing the commitment of preexisting users [Hristova et al. 2013]. Social trading, based on text-based chats among the traders in a company, has also been introduced recently [Saavedra et al. 2011]. Both of these cases are examples of machine mediated

collective behaviors in which social interactions are not included in the original design.

Efforts toward strengthening social interactions have been evident in numerous other research areas as well. For instance, Marcin et al. [2014] discuss two studies conducted with an affective dialog system in which text-based system-user communication was used to model, generate, and present different affective and social interaction scenarios in order to intensify human-human interactions within HMNs. In the same framework, the SoCS project aimed to support higher quality online deliberation in HMNs, especially by supporting a number of “social deliberative skills” such as perspective taking, empathy, self-reflection, tolerance for uncertainty, listening and question-asking skills, and meta-dialog in online contexts [Murray et al. 2014].

Second, and related to the first trend, human-machine interactions are becoming less demanding. By introducing social functionality, HMN designers capitalize on social motivations to encourage participation. In some cases, HMN designers have also started introducing monetary incentives, as in paying for contributions on review and rating sites or for crowdsensing applications. Last but not least, automated data collection is gaining prominence. These changes serve to reduce the effort on behalf of human participants and increase their gains from participation. HMNs no longer need to rely on “altruists” to continue functioning and growing. Previous research has shown that the introduction of social functionality in HMNs eases their overall operation and node interaction. For instance, Murray-Rust et al. [2014] present models and techniques for coordination of human workers in crowdsourced software development environments. The techniques augment the existing Social Compute Unit concept—a general framework for management of ad-hoc human worker teams. This approach allows the researchers to combine coordination and quality constraints with dynamic assessments of software-users’ desires, while dynamically choosing appropriate software development coordination models, leading to easier human-machine interactions.

Third, machine-human interactions are also becoming more prominent. Machines now customize and filter any information that users receive and consume. This is evident in numerous HMNs nowadays. A typical example is Facebook and the content control carried out within this network. Even in such a massive network, advanced content control algorithms are applied and have proven to be quite effective [Bakshy et al. 2015]. Similar algorithms are applied in several other HMNs enhancing machine-human interactions.

Fourth, not just the nature of interactions but also the actual structure of the HMNs is starting to evolve. There is a trend toward redesigning HMNs as peer-to-peer networks. A typical example regarding that trend is Bitcoin. The Bitcoin cryptocurrency is built on top of a decentralized P2P network used to propagate system information such as transactions or blockchain updates. The Bitcoin architecture does not rely on a centralized server. Instead, a distributed approach has been adopted to support the system [Donet Donet et al. 2014]. Such an approach can be used in many of a system’s facets, including data storage, data confirmation, and data transmission. The major reason for this is that P2P architectures improve the scalability of HMNs. But they also entail certain undesirable consequences. In P2P networks, machines become equivalent with users, as each user contributes a machine to store, process, and transmit data. This means that machines start to exhibit some of the problems that humans have: unavailability, unreliability, distrustfulness, and untrustworthiness. This implies that social science knowledge and approaches are starting to be indispensable even for the engineering of HMNs.

The fifth trend is a product of the previous four—the number of HMNs that cut across the eight types we have identified is starting to increase. For example, cryptocurrencies such as Bitcoin have the properties of both public-resource computing and

online markets. Online and mobile dating services such as Match.com and Tinder have characteristics from both online markets and social media. Wildlife@Home, a project aimed at analyzing videos of wildlife, combines public resource computing with crowdsourcing [Desell et al. 2013]. Open innovation projects often extend the crowdsourcing approach to mass collaboration, by combining ideation with collaborative refinement of ideas. Opportunistic Internet of Things systems combine crowdsensing with social media [Guo et al. 2015]. Since these hybrid types of HMNs are relatively new, the literature on them is still growing. Our survey methodology placed an emphasis on publication impact and as a result, down-weighted these networks. Nevertheless, this trend of hybridization renders our survey particularly valuable, as it emphasizes the need for the transfer of knowledge and design solutions across different types of HMNs.

4.2. Common Challenges

In this survey, we analyzed different types of interactions to identify different types of HMNs. However, there are risks and challenges in relation to the interactions that are common among all HMNs. Here we discuss issues related to user motivation, trust, user experience, privacy, and scalability.

Attracting users and keeping them motivated to participate and contribute is a major challenge across all identified HMNs. A number of approaches to mitigating this challenge have been proposed, as we have seen in Section 3. A basic distinction between these approaches concerns whether to exploit extrinsic or intrinsic motivational factors, and different types of HMNs exploit such motivational factors differently. HMNs for public resource computing and multiplayer online games typically aim for intrinsic participant motivation, for example by targeting the civic benefit or experiential value of the HMN. HMNs for crowdsensing and online markets, on the other hand, typically address recruitment and motivation challenges through extrinsic motivational factors such as remuneration or utilitarian gains. Possibly, when designing for increased user recruitment and motivation for one HMN type, it could be beneficial to look to other HMN types for inspiration. For instance, design for participant recruitment and motivation in crowdsensing systems could perhaps to a greater degree accentuate the civic value of participation and, hence, to a greater degree benefit from intrinsically motivated participation. To make such intrinsically motivated participation practically feasible, however, participation costs would need to be reduced. One way is to follow the example of crowdsourcing initiatives such as reCAPTCHA where the contributions are made as integrated parts of users' everyday behavior [von Ahn et al. 2008]. Another example of making data collection an integrated part of users' everyday behavior is the use of mobile phone network position data to monitor traffic flow [Calabrese et al. 2011]. Such integrated approaches to data collection could possibly also be applied in crowdsensing.

Other approaches to participation, recruitment, and motivation are more specific to particular HMN types. For example, the use of informal joining scripts [Von Krogh et al. 2003] has been shown to be an efficient approach of relevance to establish participant motivation in systems for mass collaboration, in particular as existing collaborators need to filter out nonserious newcomers. However, for other HMN types, such as online markets or crowdsourcing, a more adequate approach for the same purpose may be to apply traditional recommender systems [Schafer et al. 2001], as participants here typically connect in multiple brief engagements and informal joining scripts are therefore neither effective nor efficient.

Trust is another major challenge for HMNs. Some participants may struggle with a concern that the technology will undermine their own position or may be unwilling to rely entirely on the capabilities of the machines [Lee and Moray 1992; Lee and See 2004]. Others may be mainly concerned with the security of their private data [Dwyer

et al. 2007]. In response, machines need to be transparent and predictable, to create a perception of trustworthiness [Cheshire 2011; Corritore et al. 2003; McEvily et al. 2003; Schoorman et al. 2007]. On the other hand, they may also need to identify and counteract deliberate malevolence: hacking, vandalism, and intentional abuse of their algorithms, such as Google bombs and link farms.

Machine interventions to the content that observers see and use can lead to undesirable self-reinforcing feedback loops that can negatively impact the service rendered. For example, if the number of users who see a particular contribution increases the chance of others seeing it, the diversity of content can decrease. From the participants' point of view, there is again the problem of trust, since the machines in HMNs often employ "black-box" algorithms that can affect one's behavior and mood without the participant's knowledge and explicit consent [Bond et al. 2012; Kramer et al. 2014]. These two related effects have been labeled "filter bubble" [Pariser 2011]. However, such filtering or adaptation of content may be both necessary and beneficial. Without such filtering, the result may be information overload on the humans in the HMN. Further, from a machine network perspective, the size and amount of content to be delivered will impact network performance and therefore potentially undermine the Quality of Experience (QoE) for the user. This may lead to the monetization of content delivery, imposing strict controls and monitoring of throughput especially in last-mile segments.

A concern for QoE in machine-human interactions is closely related to Quality of Service (QoS) between machines within a network supporting HMN capabilities. In traditional Plain Old Telephone Service (POTS) networks, this often came down to routing and providing redundant configurations to ensure that point-to-point communication could be guaranteed irrespective of potential failures or bottlenecks. HMNs similarly require complex assemblages of multiple computers and devices; in some cases dedicated links and subnetworks are required to guarantee QoS at least (such as Content Delivery Networks). The machines communicate via Bluetooth, the Internet, or satellite, depending on proximity and the required QoS. Since not all such machines will be owned by the same service provider or carrier, contractual arrangements must be in place in support of agreed levels of service. As and when these are in place, there is also potentially the issue of transmission speeds and throughputs, leading to dedicated networks, packet-inspection, and/or local caching. Although justified on the basis of service delivery, this could compromise fairness as well as privacy [Gedik and Liu 2008; Julian et al. 2002; Kamble and Sharma 2004]. With privacy and data sensitivity in mind, there may be a need to impose a logical network on top of the existing structures (Virtual Private Network), requiring specific security protocols agreed between the different components of a network. Security architectures often mimic interpersonal trust relationships in this way, effectively introducing trust transfer into a purely machine-based exchange [Hardjono and Weis 2004; McEvily et al. 2003; Stewart 2003]. In other cases, often in connection with issues of robustness as well as privacy, dedicated networks, including private and reserved broadcast wavelengths will be used to carry data for specific purposes or in special cases [Chiti et al. 2008; Salkintzis 2006]. There may even be specific automated practices introduced to impose security compliance without human intervention.

Another significant challenge faced by most HMNs is scalability [He et al. 2011; Ugander and Backstrom 2013]. As the number of humans in the HMN grows, machines need to connect, communicate, and compute in a more efficient way, including sometimes sophisticated routing and resource allocation management to ensure connection between the right user(s) and the right machine(s) [Baran and Wu 1989; Broch et al. 1998; Cardellini et al. 1999]. In addition, they need to be able to provide consistently good bandwidth, guarantee compatibility between different communication

protocols and operating systems, and improve data structures for efficient querying, caching, and storage [Chen et al. 2008; Mehlhorn and Näher 1989].

4.3. Design Implications

Our framework helps identify common design approaches to the challenges that different HMNs share. We can group the design implications according to the four types of interactions we investigated: human-human, human input to machine, machine output to humans, and machine-machine.

Regarding human-human interactions, the major design concern is how to recruit human actors and encourage productive or constructive interaction between them. To secure a solid base of users, special attention should be given to attracting new users through peer influence and protecting them from bouncing by streamlining their first contributions and preventing harassment by more experienced users. Further, design solutions should be geared toward improving engagement and collaboration between users and creating a feeling of community. This could be achieved by, for example, organizing offline social events, symbolically rewarding users with badges, applying a loyalty ladder, strengthening within-platform communication, and strengthening trust through rich profiles and recommendations.

The design implications concerning human input are more typically associated with the User Interface (UI) and the more traditional Human-Computer Interaction (HCI) domain. They relate to encouraging as well as facilitating human interaction. Human actors need to be encouraged to use the interface and to behave in ways that provide the types of input required. This may be associated with the persuasiveness of the interface, although this is not the whole story. Continued use requires efficiency and perceived usefulness, as captured in the technology acceptance model [Venkatesh 2000]. Further, trust and privacy should be ensured, for example, by offering strict and clear privacy policies and employing transparent algorithms. Additional design solutions include organizing contributions around campaigns rather than routine as is often done with crowdsourcing projects related to natural disasters and encouraging shared responsibility for the HMN by requesting regular feedback from users on the overall state of the project.

The design issues regarding machine output concern how machines respond to humans and what they do with the information and content provided. It is important that machines respond proactively and are perceived as positive contributors. Design solutions should consider not only advanced content filtering functionality to address information overload but also algorithms to detect original and emerging content beyond popular authoritative sources. Algorithms can route content and tasks to users according to their interests and expertise in order to improve the quality of contributions and retain contributors. Further, algorithms can sift through content to select high-quality contributions to show as examples to users in order to improve their future contributions. Machines can also creatively employ the geographical distribution of the HMN users to encourage the formation of local collaboration groups, distribute work load or even server load [Ugander and Backstrom 2013], and tackle projects and tasks that are inconceivable otherwise. For example, projects that require continuous input can be intelligently routed to users from different time zones.

Regarding machine-machine interactions, it is not enough for human actors to perceive the effectiveness of the machine components they are most clearly involved with, that is, their own computers. Instead, for the HMN to be a success, there must also be a correspondingly productive interplay between the machines themselves as effective participants. Security is of primary concern here. When machine agents in a network wish to interact with others, often unknown to them, there needs to be a separate security negotiation including authentication and authorization. One way around this

would be to introduce a single “authority” who would vouch for individual machine agents to mediate their access to other services.

4.4. Future Considerations

Most previous work on networks has tended to focus either on social networks as communities of people sharing and interacting with one another online or offline, or on machine networks in terms of their robustness and complexity. It is our view, however, that in combining the two, we begin to appreciate how human-machine interactions can and often do lead to subtly interdependent behaviors that would have been impossible before. Most earlier analysis of networks and behavior has tended to treat the actors or entities within the network as equivalent within their defined type. Moreover, the characteristics and behavior of those entities were seen to remain fixed and predictable once defined [Boccaletti et al. 2006]. In order to develop a more advanced and comprehensive analysis of HMNs, we need to consider that the exhibited characteristics of nodes are dependent upon context, which may change in a complex way in response to a number of different factors [Milo 2002; Newman 2003]. For example, human actors may be described as in-group and out-group, which defines their behavior in relation to other connected humans [Giles and Giles 2012; Hogg and Reid 2006]. Similarly, a human actor may take the role of leader (or “seed”) at one point in time, but may relinquish that role in favor of other actors in authority at a later point in time [Aguirre et al. 1998; Lemonik Arthur 2013].

Machines are, as ever, more predictable than humans, but it is still conceivable that they may take more than one role in a network—for example, a mobile phone may be a communication device, a display device, or a sensor under different circumstances or at different stages of an evolving network situation [Arminen 2006; Raento et al. 2009; Vinciarelli et al. 2012]. Each entity in the network can be characterized by a number of attributes, some of which are intrinsic to the entity (such as patient, doctor, collaborator, contributor) and some of which are context sensitive (such as leader, seed, follower, in-group, out-group, passivity) [Horwitz et al. 2004; Newman 2003]. Similarly, the relationship between entities can be characterized as a function of both context and time—so trust and reliance are an attribute of the evolving relationship between two entities [Rousseau et al. 1998; Schoorman et al. 2007; Söllner et al. 2012; Söllner and Leimeister 2013], whereas trustworthiness is an attribute of a single entity fulfilling a specific role [Corritore et al. 2003; Rousseau et al. 1998]. These attributes of entities should be identified and characterized, their potential to adopt different behaviors and likelihood to adapt or respond to unexpected interactions, as well as the time-variant or dynamic nature of the relationships should be incorporated [Milo 2002; Newman 2003].

The connections between entities in a network, edges, are often taken simply to represent the existence of a relationship, effectively taking the form of an infinite capacity conduit [Boccaletti et al. 2006]. However, the edge itself may also display characteristics that influence behavior in the network. The edge is effectively a path across which information flows between two entities, but it can influence both the type of interaction and the perception by a human entity of another machine or human node [Lusseau 2003; Newman 2003]. For example, the connection may be a slow or unreliable link, it will have a capacity limit, which determines the type of information it can carry (text, data, audio, images), and as such may influence the level of trust or reliance between nodes and may therefore affect the response that a person has to the information available to them [Lu et al. 2007; Newman and Park 2003]. Future analysis therefore needs to take into account the intrinsic characteristics of the edge as well as relate it to other network characteristics that may be influenced by it.

In reality, the emergence of new behaviors in an HMN is very difficult to predict, although the effect that new user behavior has on the operation of a network can be profound [Jones 2010; Lusseau 2003]. As an example, consider the impact of social media usage on networks of people mediated by smartphones. The shift in behavior from verbal communication (telephony) to text communication has driven the developing capabilities of smart phones and the demand for universal-wide bandwidth data connection. The inclusion of cameras of ever-increasing resolution and sophistication and biosensors in smartphones, and the ubiquity of apps has resulted in behavioral change in the users of such devices and networks. In social media networks such as Twitter and Instagram, the emergence of the use of hashtags has revolutionized the sharing and discovery of information and has led to the adoption by the network of concepts such as “trending” [Brock 2012; Doctor 2013; Fortunato 2013; Glasgow and Fink 2013]. These examples are indications of the way in which human users of any technology will adapt it to their own requirements, often beyond the expectations of the originator of the technology, and, once adopted and widespread, the technology provider will then move to incorporate the new usage into their offering. This is another aspect of the dynamic nature of a technology-based network.

That networks evolve has been established earlier, along with the reasons for their evolution. Any characterisation of a network therefore needs to embody the dynamic nature of the entities and connections between them, particularly in relation to the changing context of both the individuals and the network as a whole [Fortunato 2010; Milo 2002]. Moreover, the underlying capabilities of the network will evolve with time, possibly leading to “tipping points” at which new or emergent behavior becomes possible, and this is paralleled by the developing competence of the network users whereby they are able to use the technology more effectively [Boyatzis 2006; Rogers 2008]. The very existence of the network and the sharing of experience between people using the network ensure that behavior can become “viral” and spread rapidly between users [Bakshy et al. 2012; Kietzmann et al. 2012]. This is not just “herd behavior”, although it can lead to such phenomena [Langley et al. 2014; Ronson 2015], but can also be part of a shared learning phenomenon, which can ultimately lead to modifications in the way that the technological components of the network are implemented [Benbya and McKelvey 2006].

5. CONCLUSION

This survey presented the state of the art in the field of HMNs. We focused on the actors in these networks and the interactions between them. We identified eight types of HMNs: public resource computing, crowdsourcing, web search engines, crowdsensing, online markets, social media, multiplayer online games and virtual worlds, and mass collaboration. These types differ in the structure and intensity of the interactions among and between humans and machines. We systematically collected relevant research on these types, with an emphasis on recent and high-impact work. We reviewed this work with a focus on issues related to designing the HMNs: motivating participants, guaranteeing their privacy, designing reputation, recommendation, and content-ranking algorithms, aggregating and processing contributions, and so on.

We have identified and discussed five emerging trends in HMNs: human-human interactions are intensifying in HMNs; human-machine interactions are becoming less demanding; machine-human interactions are becoming more prominent; the structure of HMNs is evolving; more hybrid types of HMNs are emerging. The five development trends suggest that the differences between the eight HMN types we identified are starting to blur. Nevertheless, our analytical framework remains useful for identifying specific niches for development and innovation. Some possibilities include citizen

science projects that encourage collaboration between users, search engines that rely on users explicitly ranking sites, and social network sites that enable file sharing.

Our framework helped us identify shared challenges, such as those related to user motivation, trust, user experience, privacy, and scalability. Our framework also helped organize the design implications from these challenges. An avenue for further work is to more systematically delineate types of HMNs and characteristics relevant to describing them in order to identify transferable solutions that may benefit ICT developers when designing new HMNs. Building upon the survey presented here, Eide et al. [2016] have already taken first steps toward a more systematic typology.

Key to our definition of an HMN is the synergistic effect that can occur from the interactions between humans and machines. Synergy may result from emergent behavior, which can be challenging to predict. We have discussed several aspects of emergent behavior and network dynamics pertaining to, for example, context of use, changing role in the network, and type of connections between actors. Characterization of an HMN needs to embody the dynamic nature of the actors and the connections between them. How to characterize HMNs to allow for the prediction of emergent behavior is another interesting avenue for further work.

Our framework has its intrinsic limitations as well. We have simplified the type of interactions and limited the possible combinations. Moreover, we only considered few modes of contributions by actors, and we did not allow the actors to take different roles. Still, the simplicity of the provided framework allows for straightforward extensions in future iterations.

REFERENCES

- Mavis T. Adjei, Stephanie M. Noble, and Charles H. Noble. 2010. The influence of C2C communications in online brand communities on customer purchase behavior. *Journal of the Academy of Marketing Science* 38 (2010), 634.
- Eugene Agichtein, Carlos Castillo, Debora Donato, Aristides Gionis, and Gilad Mishne. 2008. Finding high-quality content in social media. In *Proceedings of the International Conference on Web Search and Web Data Mining (WSDM'08)*. ACM, New York, 183.
- B. E. Aguirre, Dennis Wenger, and Gabriela Vigo. 1998. A test of the emergent norm theory of collective behavior. *Sociological Forum* 13, 2 (1998), 301–320.
- Mohammad Allahbakhsh, Boualem Benatallah, Aleksandar Ignjatovic, Hamid Reza Motahari-Nezhad, Elisa Bertino, and Schahram Dustdar. 2013. Quality control in crowdsourcing systems: Issues and directions. *IEEE Internet Computing* 17, 2 (2013), 76–81.
- David P. Anderson. 2004. BOINC: A system for public-resource computing and storage. In *Proceedings of the 5th IEEE/ACM International Workshop on Grid Computing*. IEEE, 4–10.
- David P. Anderson, Jeff Cobb, Eric Korpela, Matt Lebofsky, and Dan Werthimer. 2002. SETI@home: An experiment in public-resource computing. *Communications of the ACM* 45, 11 (2002), 56–61.
- Arvind Arasu, Junghoo Cho, Hector Garcia-Molina, Andreas Paepcke, and Sriram Raghavan. 2001. Searching the web. *ACM Transactions on Internet Technology* 1, 1 (2001), 2–43.
- Alexandre Ardichvili. 2008. Learning and knowledge sharing in virtual communities of practice: Motivators, barriers, and enablers. *Advances in Developing Human Resources* 10, 4 (2008), 541–554.
- Ilkka Arminen. 2006. Social functions of location in mobile telephony. *Personal and Ubiquitous Computing* 10, 5 (2006), 319–323.
- Luigi Atzori, Antonio Iera, and Giacomo Morabito. 2010. The internet of things: A survey. *Computer Networks* 54, 15 (2010), 2787–2805.
- Hyunmi Baek, JoongHo Ahn, and Youngseok Choi. 2012. Helpfulness of online consumer reviews: Readers' objectives and review cues. *International Journal of Electronic Commerce* 17, 2 (2012), 99–126.
- Richard P. Bagozzi and Utpal M. Dholakia. 2006. Open source software user communities: A study of participation in Linux user groups. *Management Science* 52, 7 (2006), 1099–1115.
- Eytan Bakshy, Solomon Messing, and Lada Adamic. 2015. Exposure to ideologically diverse news and opinion on facebook. *Science* 348, 6239 (2015), 1130–1132.

- Eytan Bakshy, Itamar Rosenn, Cameron Marlow, and Lada Adamic. 2012. The role of social networks in information diffusion. In *Proceedings of the 21st International Conference on World Wide Web (WWW'12)*. ACM, 519.
- Carliss Y. Baldwin and Kim B. Clark. 2006. The architecture of participation: Does code architecture mitigate free riding in the open source development model? *Management Science* 52, 7 (2006), 1116–1127.
- Mesut E. Baran and Felix F. Wu. 1989. Network reconfiguration in distribution systems for loss reduction and load balancing. *IEEE Transactions on Power Delivery* 4, 2 (1989), 1401–1407.
- Barry L. Bayus. 2013. Crowdsourcing new product ideas over time: An analysis of the Dell ideastorm community. *Management Science* 59, 1 (2013), 226–244.
- Adam L. Beberg, Daniel L. Ensign, Guha Jayachandran, Siraj Khaliq, and Vijay S. Pande. 2009. Folding@home: Lessons from eight years of volunteer distributed computing. In *Proceedings of the 2009 IEEE International Symposium on Parallel & Distributed Processing*. IEEE, 1–8.
- Hind Benbya and Bill McKelvey. 2006. Toward a complexity theory of information systems development. *Information Technology & People* 19, 1 (2006), 12–34.
- Tim Berners-Lee. 2000. *Weaving the Web: The Original Design and Ultimate Destiny of the World Wide Web*. Collins Business, London. 256 pages.
- Michael S. Bernstein, Greg Little, Robert C. Miller, Björn Hartmann, Mark S. Ackerman, David R. Karger, David Crowell, and Katrina Panovich. 2010. Soylent: A word processor with a crowd inside. In *Proceedings of the 23rd Annual ACM Symposium on User Interface Software and Technology (UIST'10)*. ACM, New York, 313.
- Christian Bird, David Pattison, Raissa D'Souza, Vladimir Filkov, and Premkumar Devanbu. 2008. Latent social structure in open source projects. In *Proceedings of the 16th ACM SIGSOFT International Symposium on Foundations of Software Engineering (SIGSOFT'08/FSE-16)*. ACM, New York, 24.
- Stefano Boccaletti, Vito Latora, Yamir Moreno, Martin Chavez, and Dong-Uk Hwang. 2006. Complex networks: Structure and dynamics. *Physics Reports* 424, 4–5 (2006), 175–308.
- Robert M. Bond, Christopher J. Fariss, Jason J. Jones, Adam D. I. Kramer, Cameron Marlow, Jaime E. Settle, and James H. Fowler. 2012. A 61-million-person experiment in social influence and political mobilization. *Nature* 489, 7415 (2012), 295–298.
- Richard E. Boyatzis. 2006. An overview of intentional change from a complexity perspective. *Journal of Management Development* 25, 7 (2006), 607–623.
- Danah M. Boyd and Nicole B. Ellison. 2007. Social network sites: Definition, history, and scholarship. *Journal of Computer-Mediated Communication* 13, 1 (2007), 210–230.
- Daren C. Brabham. 2008. Crowdsourcing as a model for problem solving. *Convergence: The International Journal of Research into New Media Technologies* 14, 1 (2008), 75–90.
- Sergey Brin and Lawrence Page. 2012. Reprint of: The anatomy of a large-scale hypertextual web search engine. *Computer Networks* 56, 18 (2012), 3825–3833.
- Josh Broch, David A. Maltz, David B. Johnson, Yih-Chun Hu, and Jorjeta Jetcheva. 1998. A performance comparison of multi-hop wireless ad hoc network routing protocols. In *Proceedings of the 4th Annual ACM/IEEE International Conference on Mobile Computing and Networking*. ACM, 85–97.
- André Brock. 2012. From the blackhand side: Twitter as a cultural conversation. *Journal of Broadcasting & Electronic Media* 56, 4 (2012), 529–549.
- Ignasi Buch, Matt J. Harvey, Toni Giorgino, David P. Anderson, and Gianni De Fabritiis. 2010. High-throughput all-atom molecular dynamics simulations using distributed computing. *Journal of Chemical Information and Modeling* 50, 3 (2010), 397–403.
- Vanilson Burégio, Silvio Meira, and Nelson Rosa. 2013. Social machines: A unified paradigm to describe social Web-oriented systems. In *International World Wide Web Conference Committee (IW3C2)*. WWW 2013 Companion, 1–5.
- Francesco Calabrese, Massimo Colonna, Piero Lovisolo, Dario Parata, and Carlo Ratti. 2011. Real-time urban monitoring using cell phones: A case study in rome. *IEEE Transactions on Intelligent Transport Systems* 12, 1 (2011), 141–151.
- Valeria Cardellini, Michele Colajanni, and S. Yu Philip. 1999. Dynamic load balancing on web-server systems. *IEEE Internet Computing* 3, 3 (1999), 28–39.
- Claudio Carpineto and Giovanni Romano. 2012. A survey of automatic query expansion in information retrieval. *ACM Computing Surveys* 44, 1 (2012), 1–50.
- Manuel Castells. 2011. *The Rise of the Network Society: The Information Age: Economy, Society, and Culture*. Vol. 1. John Wiley & Sons.
- Damon Centola and Arnout van de Rijt. 2015. Choosing your network: Social preferences in an online health community. *Social Science & Medicine (1982)* 125 (2015), 19–31.

- David Chen, Guy Doumeingts, and François Vernadat. 2008. Architectures for enterprise integration and interoperability: Past, present and future. *Computers in Industry* 59, 7 (2008), 647–659.
- Duan-Bing Chen, Hui Gao, Linyuan Lü, and Tao Zhou. 2013. Identifying influential nodes in large-scale directed networks: The role of clustering. *PloS One* 8, 10 (2013), e77455.
- Coye Cheshire. 2011. Online trust, trustworthiness, or assurance? *Daedalus* 140, 4 (2011), 49–58.
- Francesco Chiti, Romano Fantacci, Leonardo Maccari, Dania Marabissi, and Daniele Tarchi. 2008. A broadband wireless communications system for emergency management. *IEEE Wireless Communications* 15, 3 (2008), 8–14.
- Chao-Min Chiu, Meng-Hsiang Hsu, and Eric T. G. Wang. 2006. Understanding knowledge sharing in virtual communities: An integration of social capital and social cognitive theories. *Decision Support Systems* 42, 3 (2006), 1872–1888.
- Dongseong Choi and Jinwoo Kim. 2004. Why people continue to play online games: In search of critical design factors to increase customer loyalty to online contents. *Cyberpsychology & Behavior* 7, 1 (2004), 11–24.
- Delphine Christin, Andreas Reinhardt, Salil S. Kanhere, and Matthias Hollick. 2011. A survey on privacy in mobile participatory sensing applications. *Journal of Systems and Software* 84, 11 (2011), 1928–1946.
- Mark Claypool and Kaja Claypool. 2006. Latency and player actions in online games. *Communications of the ACM* 49, 11 (2006), 40.
- Kevyn Collins-Thompson, Paul N. Bennett, Ryan W. White, Sebastian de la Chica, and David Sontag. 2011. Personalizing web search results by reading level. In *Proceedings of the 20th ACM International Conference on Information and Knowledge Management (CIKM'11)*. ACM, New York, 403.
- Sunny Consolvo, Ryan Libby, Ian Smith, James A. Landay, David W. McDonald, Tammy Toscos, Mike Y. Chen, Jon Froehlich, Beverly Harrison, Predrag Klasnja, Anthony LaMarca, and Louis LeGrand. 2008. Activity sensing in the wild: A field trial of UbiFit garden. In *Proceedings of the 26th Annual CHI Conference on Human Factors in Computing Systems (CHI'08)*. ACM, New York, 1797.
- Seth Cooper, Firas Khatib, Adrien Treuille, Janos Barbero, Jeehyung Lee, Michael Beenen, Andrew Leaver-Fay, David Baker, Zoran Popović, and Foldit Players. 2010. Predicting protein structures with a multiplayer online game. *Nature* 466, 7307 (2010), 756–760.
- Cynthia L. Corritore, Beverly Kracher, and Susan Wiedenbeck. 2003. On-line trust: Concepts, evolving themes, a model. *International Journal of Human-Computer Studies* 58, 6 (2003), 737–758.
- Kevin Crowston and James Howison. 2005. The social structure of free and open source software development. *First Monday* 10, 2 (2005).
- Kevin Crowston, Kangning Wei, James Howison, and Andrea Wiggins. 2012. Free/libre open-source software development. *ACM Computing Surveys* 44, 2 (2012), 1–35.
- Antonio Cuomo, Giuseppe Di Modica, Salvatore Distefano, Antonio Puliafito, Massimiliano Rak, Orazio Tomarchio, Salvatore Venticinque, and Umberto Villano. 2012. An SLA-based broker for cloud infrastructures. *Journal of Grid Computing* 11, 1 (2012), 1–25.
- Leucio Cutillo, Refik Molva, and Thorsten Strufe. 2009. Safebook: A privacy-preserving online social network leveraging on real-life trust. *IEEE Communications Magazine* 47, 12 (2009), 94–101.
- Yves-Alexandre de Montjoye, César A. Hidalgo, Michel Verleysen, and Vincent D. Blondel. 2013. Unique in the crowd: The privacy bounds of human mobility. *Scientific Reports* 3 (2013), 1376.
- Travis Desell, Robert Bergman, Kyle Goehner, Ronald Marsh, Rebecca VanderClute, and Susan Ellis-Felege. 2013. Wildlife@Home: Combining crowd sourcing and volunteer computing to analyze avian nesting video. In *Proceedings of the 2013 IEEE 9th International Conference on e-Science*. IEEE, 107–115.
- Anhai Doan, Raghu Ramakrishnan, and Alon Y. Halevy. 2011. Crowdsourcing systems on the world-wide web. *Communications of the ACM* 54, 4 (2011), 86.
- Vanessa Doctor. 2013. Hashtag history: When and what started it? Retrieved from <https://www.hashtags.org/featured/hashtag-history-when-and-what-started-it/>.
- Judith Donath and Danah Boyd. 2004. Public displays of connection. *BT Technology Journal* 22, 4 (2004), 71–82.
- Joan Antoni Donet Donet, Cristina Pérez-Solà, and Jordi Herrera-Joancomartí. 2014. The bitcoin P2P network. In *Financial Cryptography and Data Security*, Rainer Böhme, Michael Brenner, Tyler Moore, and Matthew Smith (Eds.). Lecture Notes in Computer Science, Vol. 8438. Springer, Berlin, 87–102.
- Nicolas Ducheneaut and Robert J. Moore. 2004. The social side of gaming: A study of interaction patterns in a massively multiplayer online game. In *Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work (CSCW'04)*. ACM, New York, 360.
- Nicolas Ducheneaut, Nicholas Yee, Eric Nickell, and Robert J. Moore. 2006. Alone together? Exploring the social dynamics of massively multiplayer online games. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'06)*. ACM, New York, 407.

- Nicolas Ducheneaut, Nicholas Yee, Eric Nickell, and Robert J. Moore. 2007. The life and death of online gaming communities: A look at guilds in world of warcraft. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'07)*. ACM, New York, 839.
- Catherine Dwyer, Starr Roxanne Hiltz, and Katia Passerini. 2007. Trust and privacy concern within social networking sites: A comparison of Facebook and MySpace. Retrieved from <http://aisel.aisnet.org/cgi/viewcontent.cgi?article=1849&context=amcis2007>.
- Aslak Wegner Eide, J. Brian Pickering, Taha Yasseri, George Bravos, Asbjørn Følstad, Vegard Engen, Paul Walland, Eric T. Meyer, and Milena Tsvetkova. 2016. Human-machine networks: Towards a typology and profiling framework. In *HCI International*. Springer, Toronto, Canada.
- Fred E. Emery and Eric L. Trist. 1965. The causal texture of organizational environments. *Human Relations* 18, 1 (1965), 21–32.
- Enrique Estelles-Arolas and Fernando González-Ladrón-de Guevara. 2012. Towards an integrated crowdsourcing definition. *Journal of Information Science* 38, 2 (2012), 189–200.
- Lujun Fang and Kristen LeFevre. 2010. Privacy wizards for social networking sites. In *Proceedings of the 19th International Conference on World Wide Web (WWW'10)*. ACM, New York, 351.
- Yulin Fang and Derrick Neufeld. 2009. Understanding sustained participation in open source software projects. *Journal of Management Information Systems* 25, 4 (2009), 9–50.
- Matthew Faulkner, MingHei Cheng, Andreas Krause, Robert Clayton, Thomas Heaton, K. Mani Chandy, Monica Kohler, Julian Bunn, Richard Guy, Annie Liu, and Michael Olson. 2014. Community sense and response systems: Your phone as quake detector. *Communications of the ACM* 57, 7 (2014), 66–75.
- Joe Fortunato. 2013. The hashtag: A history deeper than Twitter. (2013). <http://www.copyright.com/blog/the-hashtag-a-history-deeper-than-twitter/>.
- Santo Fortunato. 2010. Community detection in graphs. *Physics Reports* 486, 3–5 (2010), 75–174.
- Ian Foster, Carl Kesselman, and Steven Tuecke. 2001. The anatomy of the grid: Enabling scalable virtual organizations. *International Journal of High Performance Computing Applications* 15, 3 (2001), 200–222.
- Bruno S. Frey and Reto Jegen. 1999. *Motivation Crowding Theory: A Survey of Empirical Evidence*. Institute for Empirical Research in Economics, University of Zurich.
- Jon Froehlich, Tawanna Dillahunt, Predrag Klasnja, Jennifer Mankoff, Sunny Consolvo, Beverly Harrison, and James A. Landay. 2009. UbiGreen: Investigating a mobile tool for tracking and supporting green transportation habits. In *Proceedings of the 27th International Conference on Human Factors in Computing Systems (CHI'09)*. ACM, New York, 1043.
- Raghu Ganti, Fan Ye, and Hui Lei. 2011. Mobile crowdsensing: Current state and future challenges. *IEEE Communications Magazine* 49, 11 (2011), 32–39.
- Raghu K. Ganti, Nam Pham, Hossein Ahmadi, Saurabh Nangia, and Tarek F. Abdelzaher. 2010. GreenGPS: A participatory sensing fuel-efficient maps application. In *Proceedings of the 8th International Conference on Mobile Systems, Applications, and Services (MobiSys'10)*. ACM, New York, 151.
- Bura Gedik and Ling Liu. 2008. Protecting location privacy with personalized k-anonymity: Architecture and algorithms. *IEEE Transactions on Mobile Computing* 7, 1 (2008), 1–18.
- David Geiger, Stefan Seedorf, Thimo Schulze, Robert C. Nickerson, and Martin Schader. 2011. Managing the crowd: Towards a taxonomy of crowdsourcing processes. (2011). http://aisel.aisnet.org/amcis2011_submissions/430
- Eric Gilbert and Karrie Karahalios. 2009. Predicting tie strength with social media. In *Proceedings of the 27th International Conference on Human Factors in Computing Systems (CHI'09)*. ACM, New York, 211.
- Howard Giles and Jane Giles. 2012. Ingroups and outgroups. In *Inter/Cultural Communication: Representation and Construction of Culture.*, Anastacia Kurylo (Ed.). Sage Publications, Chapter 7, 141–161.
- Kimberly Glasgow and Clayton Fink. 2013. Hashtag lifespan and social networks during the London riots. In *Social Computing, Behavioral-Cultural Modeling and Prediction*. Springer, 311–320.
- Lakshmi Goel, Norman Johnson, Iris Junglas, and Blake Ives. 2013. Predicting users' return to virtual worlds: A social perspective. *Information Systems Journal* 23, 1 (2013), 35–63.
- Sreenivas Gollapudi and Aneesh Sharma. 2009. An axiomatic approach for result diversification. In *Proceedings of the 18th International Conference on World Wide Web (WWW'09)*. ACM, New York, 381.
- Rajdeep Grewal, Gary L. Lilien, and Girish Mallapragada. 2006. Location, location, location: How network embeddedness affects project success in open source systems. *Management Science* 52, 7 (2006), 1043–1056.
- Bin Guo, Zhu Wang, Zhiwen Yu, Yu Wang, Neil Y. Yen, Runhe Huang, and Xingshe Zhou. 2015. Mobile crowd sensing and computing. *ACM Computing Surveys* 48, 1 (2015), 1–31.
- Bin Guo, Daqing Zhang, Zhu Wang, Zhiwen Yu, and Xingshe Zhou. 2013. Opportunistic IoT: Exploring the harmonious interaction between human and the internet of things. *Journal of Network and Computer Applications* 36, 6 (2013), 1531–1539.

- Johannes Hachmann, Roberto Olivares-Amaya, Sule Atahan-Evrenk, Carlos Amador-Bedolla, Roel S. Sanchez-Carrera, Aryeh Gold-Parker, Leslie Vogt, Anna M. Brockway, and Alan Aspuru-Guzik. 2011. The Harvard clean energy project: Large-scale computational screening and design of organic photovoltaics on the World Community Grid. *The Journal of Physical Chemistry Letters* 2, 17 (2011), 2241–2251.
- Jungpil Hahn, Jae Yun Moon, and Chen Zhang. 2008. Emergence of new project teams from open source software developer networks: Impact of prior collaboration ties. *Information Systems Research* 19, 3 (2008), 369–391.
- Robin Hanson, Ryan Oprea, and David Porter. 2006. Information aggregation and manipulation in an experimental market. *Journal of Economic Behavior & Organization* 60, 4 (2006), 449–459.
- Thomas Hardjono and Brian Weis. 2004. The multicast group security architecture. (2004). <https://tools.ietf.org/html/rfc3740>
- Alexander Hars and Shaosong Ou. 2001. Working for free? Motivations of participating in open source projects. In *Proceedings of the 34th Annual Hawaii International Conference on System Sciences*. IEEE.
- Yongqiang He, Rubao Lee, Yin Huai, Zheng Shao, Namit Jain, Xiaodong Zhang, and Zhiwei Xu. 2011. RCFfile: A fast and space-efficient data placement structure in MapReduce-based warehouse systems. In *Proceedings of the 27th International Conference on Data Engineering*. IEEE, 1199–1208.
- Guido Hertel, Sven Niedner, and Stefanie Herrmann. 2003. Motivation of software developers in open source projects: An internet-based survey of contributors to the linux kernel. *Research Policy* 32, 7 (2003), 1159–1177.
- Paul Heymann, Georgia Koutrika, and Hector Garcia-Molina. 2008. Can social bookmarking improve web search?. In *Proceedings of the International Conference on Web Search and Web Data Mining (WSDM'08)*. ACM, New York, 195.
- Michael A. Hogg and Scott A. Reid. 2006. Social identity, self-categorization, and the communication of group norms. *Communication Theory* 16 (2006), 7–30.
- Susan Horwitz, Thomas Reps, and David Binkley. 2004. Interprocedural slicing using dependence graphs. *ACM SIGPLAN Notices* 39, 4 (2004), 229.
- Jeff Howe. 2006. The rise of crowdsourcing. *Wired Magazine* 14, 6 (2006), 1–4.
- Desislava Hristova, Giovanni Quattrone, Afra J. Mashhadi, and Licia Capra. 2013. The life of the party: Impact of social mapping in OpenStreetMap. In *Proceedings of the 7th International Conference on Weblogs and Social Media (ICWSM'13)*. 234–243.
- Shun-Yun Hu, Jui-Fa Chen, and Tsu-Han Chen. 2006. VON: A scalable peer-to-peer network for virtual environments. *IEEE Network* 20, 4 (2006), 22–31.
- Jeff Huang and Efthimis N. Efthimiadis. 2009. Analyzing and evaluating query reformulation strategies in web search logs. In *Proceedings of the 18th ACM Conference on Information and Knowledge Management (CIKM'09)*. ACM, New York, 77.
- Kuan Lun Huang, Salil S. Kanhere, and Wen Hu. 2010. Preserving privacy in participatory sensing systems. *Computer Communications* 33, 11 (2010), 1266–1280.
- Markus C. Huebscher and Julie A. McCann. 2008. A survey of autonomic computing degrees, models, and applications. *ACM Computing Surveys* 40, 3 (2008), 7.
- Daniel Hughes, Geoff Coulson, and James Walkerdine. 2005. Free riding on gnutella revisited: The bell tolls? *IEEE Distributed Systems Online* 6, 6 (2005), 1.
- Panagiotis G. Ipeirotis, Foster Provost, and Jing Wang. 2010. Quality management on amazon mechanical turk. In *Proceedings of the ACM SIGKDD Workshop on Human Computation (HCOMP'10)*. ACM, New York, 64.
- Bernard J. Jansen, Danielle L. Booth, and Amanda Spink. 2008. Determining the informational, navigational, and transactional intent of web queries. *Information Processing & Management* 44, 3 (2008), 1251–1266.
- Haiyan Jia, Mu Wu, Eunhwa Jung, Alice Shapiro, and S. Shyam Sundar. 2012. Balancing human agency and object agency: An end-user interview study of the internet of things. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*. ACM, 1185–1188.
- Jeff Jones. 2010. The emergence and dynamical evolution of complex transport networks from simple low-level behaviours. *International Journal of Unconventional Computing* 6, 2 (2010), 125–144.
- Kiku Jones and Lori N. K. Leonard. 2008. Trust in consumer-to-consumer electronic commerce. *Information & Management* 45, 2 (2008), 88–95.
- Audun Jøsang, Roslan Ismail, and Colin Boyd. 2007. A survey of trust and reputation systems for online service provision. *Decision Support Systems* 43, 2 (2007), 618–644.
- David Julian, Mung Chiang, Daniel O'Neill, and Stephen Boyd. 2002. QoS and fairness constrained convex optimization of resource allocation for wireless cellular and ad hoc networks. In *Proceedings of the 21st*

- Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM'02)*, Vol. 2. IEEE, 477–486.
- Vihang Kamble and Vinod Sharma. 2004. A simple approach to provide QoS and fairness in internet. In *Proceedings of the International Conference on Signal Processing and Communications*. 412–417.
- Sepandar D. Kamvar, Mario T. Schlosser, and Hector Garcia-Molina. 2003. The eigentrust algorithm for reputation management in P2P networks. In *Proceedings of the 12th International Conference on World Wide Web (WWW'03)*. ACM, New York, 640.
- Andreas M. Kaplan and Michael Haenlein. 2010. Users of the world, unite! The challenges and opportunities of social media. *Business Horizons* 53, 1 (2010), 59–68.
- Alexander Kawrykow, Gary Roumanis, Alfred Kam, Daniel Kwak, Clarence Leung, Chu Wu, Eleyine Zarour, Luis Sarmenta, Mathieu Blanchette, and Jérôme Waldispühl. 2012. Phylo: A citizen science approach for improving multiple sequence alignment. *PLoS One* 7, 3 (2012), e31362.
- Jeffrey O. Kephart and David M. Chess. 2003. The vision of autonomic computing. *Computer* 36, 1 (2003), 41–50.
- Wazir Zada Khan, Yang Xiang, Mohammed Y. Aalsalem, and Quratulain Arshad. 2013. Mobile phone sensing systems: A survey. *IEEE Communications Surveys & Tutorials* 15, 1 (2013), 402–427.
- Firas Khatib, Seth Cooper, Michael D. Tyka, Kefan Xu, Ilya Makedon, Zoran Popovic, David Baker, and Foldit Players. 2011. Algorithm discovery by protein folding game players. *Proceedings of the National Academy of Sciences* 108, 47 (2011), 18949–18953.
- Jan H. Kietzmann, Bruno S. Silvestre, Ian P. McCarthy, and Leyland F. Pitt. 2012. Unpacking the social media phenomenon: Towards a research agenda. *Journal of Public Affairs* 12, 2 (2012), 109–119.
- Won Kim, Ok-Ran Jeong, and Sang-Won Lee. 2010. On social web sites. *Information Systems* 35, 2 (2010), 215–236.
- Aniket Kittur, Jeffrey V. Nickerson, Michael Bernstein, Elizabeth Gerber, Aaron Shaw, John Zimmerman, Matt Lease, and John Horton. 2013. The future of crowd work. In *Proceedings of the 2013 Conference on Computer Supported Cooperative Work (CSCW'13)*. ACM, New York, 1301.
- Aniket Kittur, Boris Smus, Susheel Khamkar, and Robert E. Kraut. 2011. CrowdForge: Crowdsourcing complex work. In *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology (UIST'11)*. ACM, New York, 43.
- Jon M. Kleinberg. 1999. Authoritative sources in a hyperlinked environment. *Journal of the ACM* 46, 5 (1999), 604–632.
- Michał Kosinski, David Stillwell, and Thore Graepel. 2013. Private traits and attributes are predictable from digital records of human behavior. *Proceedings of the National Academy of Sciences* 110, 15 (2013), 5802–5805.
- Iordanis Koutsopoulos. 2013. Optimal incentive-driven design of participatory sensing systems. In *Proceedings of the Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM'13)*. 1402–1410.
- Adam D. I. Kramer, Jamie E. Guillory, and Jeffrey T. Hancock. 2014. Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences* 111, 24 (2014), 8788–8790.
- Hanna Krasnova, Sarah Spiekermann, Ksenia Koroleva, and Thomas Hildebrand. 2010. Online social networks: Why we disclose. *Journal of Information Technology* 25, 2 (2010), 109–125.
- Haewoon Kwak, Changhyun Lee, Hosung Park, and Sue Moon. 2010. What is twitter, a social network or a news media?. In *Proceedings of the 19th International Conference on World Wide Web (WWW'10)*. ACM, New York, 591.
- Karim R. Lakhani, Kevin J. Boudreau, Po-Ru Loh, Lars Backstrom, Carliss Baldwin, Eric Lonstein, Mike Lydon, Alan MacCormack, Ramy A. Arnaout, and Eva C. Guinan. 2013. Prize-based contests can provide solutions to computational biology problems. *Nature Biotechnology* 31, 2 (2013), 108–111.
- Karim R. Lakhani and Eric von Hippel. 2003. How open source software works: “Free” user-to-user assistance. *Research Policy* 32, 6 (2003), 923–943.
- Nicholas Lane, Emiliano Miluzzo, Hong Lu, Daniel Peebles, Tanzeem Choudhury, and Andrew Campbell. 2010. A survey of mobile phone sensing. *IEEE Communications Magazine* 48, 9 (2010), 140–150.
- David J. Langley, Maarten C. Hovee, J. Roland Ortt, Nico Pals, and Bob van der Vecht. 2014. Patterns of herding and their occurrence in an online setting. *Journal of Interactive Marketing* 28, 1 (2014), 16–25.
- Bruno Latour. 2005. *Reassembling the Social—An Introduction to Actor-Network-Theory*. Oxford University Press, Oxford, UK.
- John Law. 1992. Notes on the theory of the actor-network: Ordering, strategy, and heterogeneity. *Systems Practice* 5, 4 (1992), 379–393.

- Jeehyung Lee, Wipapat Kladwang, Minjae Lee, Daniel Cantu, Martin Azizyan, Hanjoo Kim, Alex Limpaecher, Sungroh Yoon, Adrien Treuille, and Rhiju Das. 2014. RNA design rules from a massive open laboratory. *Proceedings of the National Academy of Sciences* 111, 6 (2014), 2122–2127.
- John D. Lee and Neville Moray. 1992. Trust, control strategies and allocation of function in human-machine systems. *Ergonomics* 35, 10 (1992), 1243–1270.
- John D. Lee and Katrina A. See. 2004. Trust in automation: Designing for appropriate reliance. *The Journal of the Human Factors and Ergonomics of Society* 46, 1 (2004), 50–80.
- Juung-Sik Lee and Baik Hoh. 2010. Sell your experiences: A market mechanism based incentive for participatory sensing. In *Proceedings of the 2010 IEEE International Conference on Pervasive Computing and Communications (PerCom)*. 60–68.
- Kyu-Min Lee, Byungjoon Min, and Kwang-Il Goh. 2015. Towards real-world complexity: An introduction to multiplex networks. *The European Physical Journal B* 88, 2 (2015), 1–20.
- Mikaila Mariel Lemonik Arthur. 2013. Emergent norm theory. In *The Wiley-Blackwell Encyclopedia of Social and Political Movements*. Blackwell Publishing Ltd, Oxford, UK.
- Paul M. Leonardi. 2012. Materiality, sociomateriality, and socio-technical systems: What do these terms mean? How are they related? Do we need them? In *Materiality and Organizing: Social Interaction in a Technological World*, Paul M. Leonardi, B. A. Nardi, and J. Kallinikos (Eds.). Oxford University Press, Oxford, 25–48.
- Loet Leydesdorff, Félix de Moya-Anegón, and Vicente P. Guerrero-Bote. 2015. Journal maps, interactive overlays, and the measurement of interdisciplinarity on the basis of scopus data (1996–2012). *Journal of the Association for Information Science and Technology* 66, 5 (2015), 1001–1016.
- Loet Leydesdorff, Ismael Rafols, and Chaomei Chen. 2013. Interactive overlays of journals and the measurement of interdisciplinarity on the basis of aggregated journal-journal citations. *Journal of the American Society of Information Science & Technology* 64, 12 (2013), 2573–2586.
- Qinghua Li, Guohong Cao, and Thomas F. La Porta. 2014. Efficient and privacy-aware data aggregation in mobile sensing. *IEEE Transactions on Dependable and Secure Computing* 11, 2 (2014), 115–129.
- Carolyn A. Lin. 2003. An interactive communication technology adoption model. *Communication Theory* 13, 4 (2003), 345–365.
- Heshan Lin, Xiaosong Ma, Jeremy Archuleta, Wu-chun Feng, Mark Gardner, and Zhe Zhang. 2010. MOON: MapReduce on opportunistic environments. In *Proceedings of the 19th ACM International Symposium on High Performance Distributed Computing (HPDC'10)*. ACM, New York, 95.
- Kuan-Yu Lin and Hsi-Peng Lu. 2011. Why people use social networking sites: An empirical study integrating network externalities and motivation theory. *Computers in Human Behavior* 27, 3 (2011), 1152–1161.
- Yingzi Lin, W. J. Zhang, and L. Glen Watson. 2003. Using eye movement parameters for evaluating human-machine interface frameworks under normal control operation and fault detection situations. *International Journal of Human-Computer Studies* 59, 6 (2003), 837–873.
- Jan Lorenz, Heiko Rauhut, Frank Schweitzer, and Dirk Helbing. 2011. How social influence can undermine the wisdom of crowd effect. *Proceedings of the National Academy of Sciences* 108, 22 (2011), 9020–9025.
- Long J. Lu, Andrea Sboner, Yuanpeng J. Huang, Hao Xin Lu, Tara A Gianoulis, Kevin Y. Yip, Philip M. Kim, Gaetano T. Montelione, and Mark B. Gerstein. 2007. Comparing classical pathways and modern networks: Towards the development of an edge ontology. *Trends in Biochemical Sciences* 32, 7 (2007), 320–331.
- Yaobin Lu, Ling Zhao, and Bin Wang. 2010. From virtual community members to C2C e-commerce buyers: Trust in virtual communities and its effect on consumers' purchase intention. *Electronic Commerce Research and Applications* 9, 4 (2010), 346–360.
- David Lusseau. 2003. The emergent properties of a dolphin social network. *Proceedings of the Royal Society B: Biological Sciences* 270, Suppl._2 (2003), S186–S188.
- Fabrizio Marozzo, Domenico Talia, and Paolo Trunfio. 2012. P2P-MapReduce: Parallel data processing in dynamic cloud environments. *Journal of Computer and System Sciences* 78, 5 (2012), 1382–1402.
- Winter Mason and Duncan J. Watts. 2010. Financial incentives and the “performance of crowds.” *ACM SIGKDD Explorations Newsletter* 11, 2 (2010), 100.
- Philip Mayer, José Velasco, Annabelle Klarl, Rolf Hennicker, Mariachiara Puviani, Francesco Tiezzi, Rosario Pugliese, Jaroslav Kezníkl, and Tomáš Bureš. 2015. The autonomic cloud. In *Software Engineering for Collective Autonomic Systems*, Martin Wirsing, Matthias Hözl, Nora Koch, and Philip Mayer (Eds.). Lecture Notes in Computer Science, Vol. 8998. Springer International Publishing, 495–512.
- Bill McEvily, Vincenzo Perrone, and Akbar Zaheer. 2003. Trust as an organizing principle. *Organization Science* 14, 1 (2003), 91–103.

- Kurt Mehlhorn and Stefan Näher. 1989. LEDA a library of efficient data types and algorithms. In *Mathematical Foundations of Computer Science 1989*. Springer, 88–106.
- Diego Mendez, Miguel Labrador, and Kandethody Ramachandran. 2013. Data interpolation for participatory sensing systems. *Pervasive and Mobile Computing* 9, 1 (2013), 132–148.
- Ron Milo. 2002. Network motifs: Simple building blocks of complex networks. *Science* 298, 5594 (2002), 824–827.
- Alan Mislove, Massimiliano Marcon, Krishna P. Gummadi, Peter Druschel, and Bobby Bhattacharjee. 2007. Measurement and analysis of online social networks. In *Proceedings of the 7th ACM SIGCOMM Conference on Internet Measurement (IMC'07)*. ACM, New York, 29.
- Audris Mockus, Roy T. Fielding, and James D. Herbsleb. 2002. Two case studies of open source software development: Apache and Mozilla. *ACM Transactions on Software Engineering and Methodology* 11, 3 (2002), 309–346.
- Meredith Ringel Morris. 2013. Collaborative search revisited. In *Proceedings of the 2013 Conference on Computer Supported Cooperative Work (CSCW'13)*. ACM, New York, 1181.
- Meredith Ringel Morris and Eric Horvitz. 2007. SearchTogether. In *Proceedings of the 20th Annual ACM Symposium on User Interface Software and Technology (UIST'07)*. ACM, New York, 3.
- Meredith Ringel Morris, Jaime Teevan, and Katrina Panovich. 2010. What do people ask their social networks, and why? In *Proceedings of the 28th International Conference on Human Factors in Computing Systems (CHI'10)*. ACM, New York, 1739.
- Lev Muchnik, Sinan Aral, and Sean J. Taylor. 2013. Social influence bias: A randomized experiment. *Science (New York, N.Y.)* 341, 6146 (2013), 647–651.
- Tom Murray, Beverly Woolf, Ethan Katsh, Leon Osterweil, Lori Carlke, and Leah Wing. 2014. SoCS: The fourth party: Improving computer-mediated deliberation through cognitive, social and emotional support. In *Final Report of the Social Deliverative Skills (SoCS) Project*. ACM, 62.
- Dave Murray-Rust, Ognjen Scekcic, Hong-Linh Truong, Dave Robertson, and Schahram Dustdar. 2014. A collaboration model for community-based software development with social machines. In *Proceedings of the COLLABORATECOM Conference*. ICST, 10.
- Mark E. J. Newman. 2003. The structure and function of complex networks. *SIAM Review* 45, 2 (2003), 167–256.
- Mark E. J. Newman and Juyong Park. 2003. Why social networks are different from other types of networks. *Physical Review E* 68, 3 (2003), 036122.
- Donald A. Norman, Andrew Ortony, and Daniel M. Russell. 2003. Affect and machine design: Lessons for the development of autonomous machines. *IBM Systems Journal* 42, 1 (2003), 38–44.
- Oded Nov, David Anderson, and Ofer Arazy. 2010. Volunteer computing: A model of the factors determining contribution to community-based scientific research. In *Proceedings of the 19th International Conference on World Wide Web (WWW'10)*. ACM, New York, 741.
- Cory R. Ondrejka. 2004. Aviators, moguls, fashionistas and barons: Economics and ownership in second life. Available at SSRN: <https://ssrn.com/abstract=614663> (2004).
- Jukka-Pekka Onnela and Felix Reed-Tsochas. 2010. Spontaneous emergence of social influence in online systems. *Proceedings of the National Academy of Sciences* 107, 43 (2010), 18375–18380.
- Shaul Oreg and Oded Nov. 2008. Exploring motivations for contributing to open source initiatives: The roles of contribution context and personal values. *Computers in Human Behavior* 24, 5 (2008), 2055–2073.
- Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. 1999. The PageRank citation ranking: Bringing order to the Web. Retrieved from <http://ilpubs.stanford.edu:8090/422/1/1999-66.pdf>.
- Maja Pantic and Leon J. M. Rothkrantz. 2003. Toward an affect-sensitive multimodal human-computer interaction. *Proceedings of the IEEE* 91, 9 (2003), 1370–1390.
- Eli Pariser. 2011. *The Filter Bubble: What the Internet is Hiding from You*. Penguin, UK.
- Veljko Pejovic and Mirco Musolesi. 2015. Anticipatory mobile computing: A survey of the state of the art and research challenges. *ACM Computing Surveys* 47, 3 (2015), 1–29.
- Rosalind W. Picard, Elias Vyzas, and Jennifer Healey. 2001. Toward machine emotional intelligence: Analysis of affective physiological state. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 23, 10 (2001), 1175–1191.
- Johan A. Pouwelse, Pawel Garbacki, Jun Wang, Arno Bakker, Jie Yang, Alexandru Iosup, Dick H. J. Epema, Marcel Reinders, Maarten R. van Steen, and Henk J. Sips. 2008. TRIBLER: A social-based peer-to-peer system. *Concurrency and Computation: Practice and Experience* 20, 2 (2008), 127–138.
- Tobias Prokesch, Heiko A. von der Gracht, and Holger Wohlenberg. 2015. Integrating prediction market and delphi methodology into a foresight support system Insights from an online game. *Technological Forecasting and Social Change* 97 (2015), 47–64.

- Alexander J. Quinn and Benjamin B. Bederson. 2011. Human computation: A survey and taxonomy of a growing field. In *Proceedings of the 2011 Annual Conference on Human Factors in Computing Systems (CHI'11)*. ACM, New York, 1403.
- Moo-Ryong Ra, Bin Liu, Tom F. La Porta, and Ramesh Govindan. 2012. Medusa: A programming framework for crowd-sensing applications. In *Proceedings of the 10th International Conference on Mobile Systems, Applications, and Services (MobiSys'12)*. ACM, New York, 337.
- Mika Raento, Antti Oulasvirta, and Nathan Eagle. 2009. Smartphones: An emerging tool for social scientists. *Sociological Methods & Research* 37, 3 (2009), 426–454.
- Vikas C. Raykar and Shipeng Yu. 2011. Ranking annotators for crowdsourced labeling tasks. In *Advances in Neural Information Processing Systems*. 1809–1817.
- Vikas C. Raykar, Shipeng Yu, Linda H. Zhao, Gerardo Hermosillo Valadez, Charles Florin, Luca Bogoni, and Linda Moy. 2010. Learning from crowds. *The Journal of Machine Learning Research* 11 (2010), 1297–1322.
- Patricia J. Rogers. 2008. Using programme theory to evaluate complicated and complex aspects of interventions. *Evaluation* 14, 1 (2008), 29–48.
- Jon Ronson. 2015. *So You've Been Publicly Shamed*. Picador, Oxford, England.
- Jeremy Rose, Matthew Jones, and Duane Truex. 2005. Socio-theoretic accounts of IS: The problem of agency. *Scandinavian Journal of Information Systems* 17, 1 (2005), 8.
- Denise M. Rousseau, Sim B. Sitkin, Ronald S. Burt, and Colin Camerer. 1998. Not so different after all: A cross-discipline view of trust. *Academy of Management Review* 23, 3 (1998), 393–404.
- Serguei Saavedra, Kathleen Hagerty, and Brian Uzzi. 2011. Synchronicity, instant messaging, and performance among financial traders. *Proceedings of the National Academy of Sciences* 108, 13 (2011), 5296–5301.
- Apostolis K. Salkintzis. 2006. Evolving public safety communication systems by integrating WLAN and TETRA networks. *IEEE Communications Magazine* 44, 1 (2006), 38–46.
- Stefan Saroiu, P. Krishna Gummadi, and Steven D. Gribble. 2001. A measurement study of peer-to-peer file sharing systems. In *Electronic Imaging 2002*, Martin G. Kienzle and Prashant J. Shenoy (Eds.). International Society for Optics and Photonics, 156–170.
- Gregory D. Saxton, Onook Oh, and Rajiv Kishore. 2013. Rules of crowdsourcing: Models, issues, and systems of control. *Information Systems Management* 30, 1 (2013), 2–20.
- Salvatore Scellato, Anastasios Noulas, and Cecilia Mascolo. 2011. Exploiting place features in link prediction on location-based social networks. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD'11)*. ACM, New York, 1046.
- J. Ben Schafer, Joseph Konstan, and John Riedi. 1999. Recommender systems in e-commerce. In *Proceedings of the 1st ACM Conference on Electronic Commerce (EC'99)*. ACM, New York, 158–166.
- J. Ben Schafer, Joseph A. Konstan, and John Riedl. 2001. E-commerce recommender recommendation applications. In *Applications of Data Mining to Electronic Commerce*. Springer, 115–153.
- F. David Schoorman, Roger C. Mayer, and James H. Davis. 2007. An integrative model of organizational trust: Past, present, and future. *Academy of Management Review* 32, 2 (2007), 344–354.
- Nigel R. Shadbolt, Daniel Alexander Smith, Elena Simperl, Max Van Kleek, Yang Yang, and Wendy Hall. 2013. Towards a classification framework for social machines. In *Proceedings of the 22nd International Conference on World Wide Web (WWW'13)*. ACM, 905–912.
- Sonali K. Shah. 2006. Motivation, governance, and the viability of hybrid forms in open source software development. *Management Science* 52, 7 (2006), 1000–1014.
- Kim Bartel Sheehan. 2002. Toward a typology of internet users and online privacy concerns. *The Information Society* 18, 1 (2002), 21–32.
- Xiang Sheng, Jian Tang, and Weiyi Zhang. 2012. Energy-efficient collaborative sensing with mobile phones. In *Proceedings of the Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM'12)*. IEEE, 1916–1924.
- Amit Sheth, Pramod Anantharam, and Cory Henson. 2013. Physical-cyber-social computing: An early 21st century approach. *IEEE Intelligent Systems* 28, 1 (2013), 78–82.
- Jonathan Silvertown, Martin Harvey, Richard Greenwood, Mike Dodd, Jon Rosewell, Tony Rebelo, Janice Ansine, and Kevin McConway. 2015. Crowdsourcing the identification of organisms: A case-study of iSpot. *ZooKeys* 480 (2015), 125–46.
- Param Vir Singh, Yong Tan, and Vijay Mookerjee. 2011. Network effects: The influence of structural social capital on open source project success. *Management Information Systems Quarterly* 35, 4 (2011), 813–829.

- Marcin Skowron, Stefan Rank, Aleksandra Świdarska, Dennis Küster, and Arvid Kappas. 2014. Applying a text-based affective dialogue system in psychological research: Case studies on the effects of system behaviour, interaction context and social exclusion. *Cognitive Computation* 6, 4 (2014), 872–891.
- Deb Sledgianowski and Songpol Kulviwat. 2009. Using social network sites: The effects of playfulness, critical mass and trust in a hedonic context. *The Journal of Computer Information Systems* 49, 4 (2009), 74–83.
- Paul R. Smart, Elena Simperl, and Nigel Shadbolt. 2014. A taxonomic framework for social machines. In *Social Collective Intelligence*. Springer, 51–85.
- Matthias Söllner, Axel Hoffmann, Holger Hoffmann, Arno Wacker, and Jan Marco Leimeister. 2012. Understanding the formation of trust in IT artifacts. In *ICIS 2012 Proceedings*. Association for Information Systems, 1–18.
- Matthias Söllner and Jan Marco Leimeister. 2013. What we really know about antecedents of trust: A critical review of the empirical information systems literature on trust. In *Psychology of Trust: New Research*, David Gefen (Ed.). Nove Science Publishers, Hauppauge, NY, Chapter 7, 127–155.
- Kunwadee Sripanidkulchai, Bruce Maggs, and Hui Zhang. 2003. Efficient content location using interest-based locality in peer-to-peer systems. In *Proceedings of the 22nd Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM'03)*, Vol. 3. IEEE, 2166–2176.
- Stephen S. Standifird. 2001. Reputation and e-commerce: eBay auctions and the asymmetrical impact of positive and negative ratings. *Journal of Management* 27, 3 (2001), 279–295.
- Ben Steichen, Helen Ashman, and Vincent Wade. 2012. A comparative survey of personalised information retrieval and adaptive hypermedia techniques. *Information Processing & Management* 48, 4 (2012), 698–724.
- Constance A. Steinkuehler and Dmitri Williams. 2006. Where everybody knows your (screen) name: Online games as “third places.” *Journal of Computer-Mediated Communication* 11, 4 (2006), 885–909.
- Katherine J. Stewart. 2003. Trust transfer on the World Wide Web. *Organization Science* 14, 1 (2003), 5–17.
- James Surowiecki. 2005. *The Wisdom of Crowds*. Anchor, New York.
- Michael Szell, Renaud Lambiotte, and Stefan Thurner. 2010. Multirelational organization of large-scale social networks in an online world. *Proceedings of the National Academy of Sciences* 107, 31 (2010), 13636–13641.
- Don Tapscott and Anthony Williams. 2011. *Wikinomics*. Atlantic Books, Limited. 300 pages.
- Ching-I Teng. 2010. Customization, immersion satisfaction, and online gamer loyalty. *Computers in Human Behavior* 26, 6 (2010), 1547–1554.
- Catalina L. Toma and Jeffrey T. Hancock. 2013. Self-affirmation underlies Facebook use. *Personality & Social Psychology Bulletin* 39, 3 (2013), 321–331.
- Johan Ugander and Lars Backstrom. 2013. Balanced label propagation for partitioning massive graphs. In *Proceedings of the 6th ACM International Conference on Web Search and Data Mining (WSDM'13)*. ACM, New York, 507.
- Viswanath Venkatesh. 2000. Determinants of perceived ease of use: Integrating control, intrinsic motivation, and emotion into the technology acceptance model. *Information Systems Research* 11, 4 (2000), 342–365.
- Tibert Verhagen, Frans Feldberg, Bart van den Hooff, Selmar Meents, and Jani Merikivi. 2012. Understanding users' motivations to engage in virtual worlds: A multipurpose model and empirical testing. *Computers in Human Behavior* 28, 2 (2012), 484–495.
- Alessandro Vinciarelli, Maja Pantic, Dirk Heylen, Catherine Pelachaud, Isabella Poggi, Francesca D'Errico, and M. Schroeder. 2012. Bridging the gap between social animal and unsocial machine: A survey of social signal processing. *IEEE Transactions on Affective Computing* 3, 1 (2012), 69–87.
- Luis von Ahn, Benjamin Maurer, Colin McMillen, David Abraham, and Manuel Blum. 2008. reCAPTCHA: Human-based character recognition via web security measures. *Science* 321, 5895 (2008), 1465–1468.
- Georg Von Krogh, Sebastian Spaeth, and Karim R. Lakhani. 2003. Community, joining, and specialization in open source software innovation: A case study. *Research Policy* 32, 7 (2003), 1217–1241.
- Fei-Yue Wang. 2010. The emergence of intelligent enterprises: From CPS to CPSS. *IEEE Intelligent Systems* 25, 4 (2010), 85–88.
- Xinlei Wang, Wei Cheng, Prasant Mohapatra, and Tarek Abdelzaher. 2013. Artsense: Anonymous reputation and trust in participatory sensing. In *Proceedings of the Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM'13)*. IEEE, 2517–2525.
- Steve R. White, James E. Hanson, Ian Whalley, David M. Chess, and Jeffrey O. Kephart. 2004. An architectural approach to autonomic computing. In *Null*. IEEE, 2–9.
- Justin Wolfers and Eric Zitzewitz. 2004. Prediction markets. *The Journal of Economic Perspectives* 18, 2 (2004), 107–126.

- Jen-Her Wu, Shu-Ching Wang, and Ho-Huang Tsai. 2010. Falling in love with online games: The uses and gratifications perspective. *Computers in Human Behavior* 26, 6 (2010), 1862–1871.
- Chenyan Xu, Sherry Ryan, Victor Prybutok, and Chao Wen. 2012. It is not for fun: An examination of social network site usage. *Information & Management* 49, 5 (2012), 210–217.
- Amir Yahyavi and Bettina Kemme. 2013. Peer-to-peer architectures for massively multiplayer online games. *ACM Computing Surveys* 46, 1 (2013), 1–51.
- Nick Yee. 2006. Motivations for play in online games. *Cyberpsychology & Behavior* 9, 6 (2006), 772–775.
- Xi Zhang, Patricia Ordóñez de Pablos, Xiaojiong Wang, Weiguang Wang, Yongqiang Sun, and Jinghuai She. 2014. Understanding the users' continuous adoption of 3D social virtual world in China: A comparative case study. *Computers in Human Behavior* 35 (2014), 578–585.
- Elena Zheleva and Lise Getoor. 2009. To join or not to join. In *Proceedings of the 18th International Conference on World Wide Web (WWW'09)*. ACM, New York, 531.
- Pengfei Zhou, Yuanqing Zheng, and Mo Li. 2012. How long to wait?: Predicting bus arrival time with mobile phone based participatory. In *Proceedings of the 10th International Conference on Mobile Systems, Applications, and Services (MobiSys'12)*. ACM, New York, 379.
- Runfang Zhou and Kai Hwang. 2007. Powertrust: A robust and scalable reputation system for trusted peer-to-peer computing. *IEEE Transactions on Parallel and Distributed Systems* 18, 4 (2007), 460–473.
- Zhongyun Zhou, Xiao-Ling Jin, Douglas R. Vogel, Yulin Fang, and Xiaojian Chen. 2011. Individual motivations and demographic differences in social virtual world uses: An exploratory investigation in second life. *International Journal of Information Management* 31, 3 (2011), 261–271.
- Matthew Zook, Mark Graham, Taylor Shelton, and Sean Gorman. 2010. Volunteered geographic information and crowdsourcing disaster relief: A case study of the haitian earthquake. *World Medical & Health Policy* 2, 2 (2010), 6–32.

Received November 2015; revised August 2016; accepted January 2017