

Where the smart things are: social machines and the Internet of Things

Paul Smart¹  · Aastha Madaan¹ · Wendy Hall¹

Published online: 17 July 2018
© The Author(s) 2018

Abstract The emergence of large-scale social media systems, such as Wikipedia, Facebook, and Twitter, has given rise to a new multi-disciplinary effort based around the concept of *social machines*. For the most part, this research effort has limited its attention to the study of Web-based systems. It has also, perhaps unsurprisingly, tended to highlight the social scientific relevance of such systems. The present paper seeks to expand the scope of the social machine research effort to encompass the Internet of Things. One advantage of this expansion is that it helps to reveal some of the links between the science of social machines and the sciences of the mind. A second advantage is that it furthers our conceptual understanding of social machines and supports the quest to derive a philosophically-robust definition of the term “social machine.” The results of the present analysis suggest that social machines are best conceived as systems in which a combination of social and technological elements play a role in the mechanistic realization of system-level phenomena. The analysis also highlights the relevance of cognitive science and the philosophy of mind to our general understanding of systems that transcend the cyber, physical, and social domains.

Keywords Internet of things · Artificial intelligence · Ambient intelligence · Social machines · Distributed cognition · Social computation

✉ Paul Smart
ps02v@ecs.soton.ac.uk

¹ Electronics and Computer Science, University of Southampton, Highfield, Southampton SO17 1BJ, UK

1 Introduction

In recent years, advances in electronic and computational engineering have yielded a range of networked devices that are poised to extend the traditional reach of the Internet to the realm of everyday physical objects. The emergence of such devices is consistent with a technological trend that has come to be known as the Internet of Things (IoT) (Atzori et al. 2010; Miller 2015; Greengard 2015). Although there is no standard definition of the term “Internet of Things,” it is typically used to refer to states-of-affairs in which everyday objects are able to exchange information via the use of standardized (Internet-based) communication protocols (e.g., Rose et al. 2015). As a result of the IoT, we are said to be on the verge of a new era—an era that features a dizzying array of so-called *smart objects* or *smart things*, e.g., smart dust, smart TVs, smart cars, smart buildings, smart cities, smart appliances, smart clothes, and so on (see Miller 2015). What makes these things worthy of the label “smart” is typically cashed out in terms of the computational and communicative potential of IoT devices. But there is, in addition, a sense in which the source of the smarts derives from the interactions between Internet-enabled objects. At a global level, therefore, the IoT conjures up a vision of the world as itself a form of intelligent system—a ‘smart world’ driven, perhaps, by a ‘global brain’ (see Heylighen 2013).¹

The impact of the IoT is likely to be felt across all sectors of society, including a number of scientific disciplines. Although estimates vary, it is widely acknowledged that billions of objects will be connected to the Internet in coming years, with as many as 100 billion objects forming part of the IoT by 2025 (see Rose et al. 2015). As noted by a number of commentators, the scale of this expansion is likely to reverberate across society, with potentially profound implications for the way we live and work (Walport 2014; Rose et al. 2015). This, then, reveals the relevance of the IoT to disciplines that are commonly grouped under the rubric of the “social sciences” (i.e., the disciplines of sociology, economics, politics, law, and so on).

But it is not just the social sciences that stand to be affected by the IoT; the IoT is also relevant to disciplines that are sometimes grouped under the general heading of the “sciences of the mind.” These include the disciplines of cognitive science, neuroscience, artificial intelligence (AI), machine learning, and the philosophy of mind. Some insight into the nature of this relevance will be revealed during the course of the present paper; but, for present purposes, it is worth noting that the IoT is often seen to provide the technological foundation for so-called ambient intelligence (AmI) (or assistive intelligence) systems (Aarts and Encarnação 2008), i.e., systems that seek to exploit the IoT for the purposes of enhancing or augmenting human cognition (see Section 2).

¹The historical precursors of such a vision date back to the earliest days of wireless technology. In 1926, for example, Nikola Tesla is purported to have said that “[w]hen wireless is perfectly applied, the whole earth will be converted into a huge brain, which in fact it is, all things being particles of a real and rhythmic whole” (see Kennedy 1926).

One approach to understanding the social and cognitive significance of the IoT is via the notion of *social machines* (Berners-Lee and Fischetti 1999; Shadbolt et al. 2016; Palermos 2017; Smart and Shadbolt 2014; Hooper et al. 2016; Hendler and Berners-Lee 2010). The term “social machine” denotes a research effort that has, for the most part, been situated within the larger disciplinary context of Web science (see Berners-Lee et al. 2006; Hendler et al. 2008). Berners-Lee and Fischetti (1999) were among the first to discuss social machines in a Web-based context. They defined a social machine as a system that features a division of labor (or division of responsibility) between the social and technological elements of a Web-based system. In particular, Berners-Lee and Fischetti (1999) suggested that we should think of social machines as systems in which the human (social) elements perform something of a ‘creative’ role, while the technological elements are involved in more ‘administrative’ functions. Since then, a variety of alternative proposals have emerged regarding the proper way to conceptualize the term “social machine.” Hooper et al. (2016), for example, conceive of social machines as problem-solving organizations, i.e., socio-technical systems that aim to resolve problems in a socially-distributed fashion. A not altogether incompatible view is proposed by Palermos (2017). Palermos suggests that we should see social machines as a form of distributed cognitive system (see Hutchins 1995), i.e., as a system in which cognitive processing routines (e.g., problem-solving processes) are realized by an interacting nexus of material elements that includes both human agents and technological artifacts.

Although a commonly-accepted definition of the term “social machine” remains elusive, there is widespread agreement about the key features of social machines. For the most part, social machines are thus viewed as systems that involve both social participation and technological mediation. In this sense, the Internet and the Web are the perfect places to look for social machines. Given their status as global information and communication networks that rely on the use of standardized communication protocols, the Internet and Web are ideally suited to supporting the kinds of (typically large-scale) social participation that pique the interests of the social machine community. Popular targets of the social machine research effort thus include social networking systems (e.g., Facebook), microblogging services (e.g., Twitter), collaborative authoring environments (e.g., Wikipedia), media sharing sites (e.g., YouTube), and citizen science platforms (e.g., Galaxy Zoo). These systems, it should be clear, have transformed the way a number of social activities are performed, and this (at least in part) highlights the value and importance of the social machine research effort. The general aim of the social machine research effort can thus be (broadly) construed as the attempt to understand the impact of technology-mediated forms of social participation, both for ourselves (as individuals) and for the societies in which we live.

The present paper attempts to advance the philosophical debate associated with the emerging science of social machines. In particular, we attempt to address the following issues:

1. Firstly, we aim to highlight the relevance of the IoT to the science of social machines. While previous research has tended to limit its attention to Web-based

- systems (e.g., Twitter, Wikipedia, and Galaxy Zoo), we attempt to show why the empirical targets of the social machine research effort should be expanded to include the IoT.
2. Secondly, we aim to highlight the interdisciplinary nature of the social machine research effort. In particular, we aim to draw attention to the conceptual linkages that exist between the science of social machines and the sciences of the mind (most notably with the disciplines of cognitive science, philosophy of mind, machine learning, and AI).
 3. Thirdly, we aim to advance our conceptual understanding of social machines by showing how a consideration of the IoT leads to new ways of viewing (and perhaps defining) social machines. The present paper thus forms part of what might be called the *definitional endeavor* associated with the social machine research effort. Our goal in this respect is not to discredit or displace existing theoretical accounts of social machines. Rather, we aim to progress the philosophical debate in the direction of an overarching theoretical position—one that appeals to a number of ostensibly distinct (and perhaps competing) constituencies.

The structure of the present paper is organized around a number of views of social machines. These views are informed by the specific properties of IoT technology, and they also help to reveal a number of important linkages between the science of social machines and the sciences of the mind. As we shall see, such points of interdisciplinary contact suggest that cognitive science and the philosophy of mind may be useful, if not indispensable, components of the social machine research effort.

2 Ambient intelligence

The IoT promises to transform the nature of our everyday interaction with physical objects, altering the way that objects are used to perform cognitive tasks and providing new opportunities to adapt the local environment to support cognitive processing. From a purely cognitive perspective, therefore, the IoT provides us with a vision of a smart world or smart environment (see Miller 2015)—an environment in which everyday objects are able to anticipate our goals, interests, needs, and concerns and adjust their behavior accordingly.

Such a vision is nicely captured by work that goes under the heading of AmI. One of the goals of AmI research is thus to create environments that work in concert with the human subject as a means of facilitating the performance of cognitive tasks. Consider, for example, the way that Ricci et al. (2015) describe the transformational impact of the IoT and associated technologies:

All this leads to a new kind of smart space in which digital, physical, and social layers are strongly intertwined. These spaces extend the classic assistive functionality of AmI toward more proactive possibilities, where the smart environment not only monitors people as they perform tasks, or supports them by executing their requests, but also influences and changes their plans and intentions. (Ricci et al. 2015, p. 60)

Although work that goes under the heading of AmI tends to restrict its attention to individual human subjects, there is no reason why the notion of AmI cannot be applied to cognitive tasks that feature the involvement of *multiple* individuals. This is important, for tasks of this sort—cognitive tasks involving the coordinated activity of multiple individuals—are often labeled as instances of distributed or collective cognition (Hutchins 1995; Hollan et al. 2000).

It is at this point that we encounter our first view of social machines. This is a view we will dub the *ambient intelligence view*:

Ambient Intelligence View

Social machines are technological systems that adapt the properties of the local environment to support the execution of collective or distributed cognitive tasks.

According to the ambient intelligence view, a social machine is a technological system whose goal is to yield a particular form of AmI; i.e., it is a system that works to ensure an appropriate adjustment of the ‘local’² environment so as to support the execution of distributed or collective cognitive processes. The *social* status of such systems is, of course, tied to the socially-distributed nature of the cognitive task that is being performed, i.e., the fact that some episode of cognitive processing is distributed across multiple human agents.

Significantly, the ambient intelligence view is one that draws on (and perhaps depends on) the kinds of tracking, monitoring, and processing capabilities that are enabled by the IoT. A social machine, under the ambient intelligence view, is thus a system that relies on an ability to adapt the local environments of multiple human individuals, an ability which itself depends on a capacity to integrate information from (possibly) remote locations and formulate decisions about how the environment of each individual should be adjusted in a manner that works to the overall cognitive good of the larger socio-technical ensemble. This, it should be clear, is a capability that is (at best) difficult to accomplish in the absence of a sufficiently rich array of networked devices, a means of integrating (or at least communicating) the information provided by those devices, and a means of coordinating device activity across the entire socio-technical organization. These demands are all ones that are easily accommodated by the IoT vision.

The ambient intelligence view of social machines is also useful when it comes to highlighting the relevance of cognitive science to the social machine research effort. The appeal to distributed cognition, for instance, establishes a relatively clear and compelling link to existing work in cognitive science—e.g., work relating to team

²The term “local” is somewhat problematic here. This is because the relevant environment is one that needs to accommodate multiple individuals. From the perspective of distributed cognition, there is no requirement for such individuals to be physically co-located—e.g., located in the same room. Instead, the notion of distributed cognition allows for situations in which individuals may be widely dispersed. There is no reason, for example, why distributed cognitive activity could not occur between the inhabitants of a future lunar colony and a collection of terrestrial (i.e., Earth-bound) collaborators.

cognition (Cooke et al. 2007), group cognition (Theiner 2014), and distributed cognition (Hutchins 1995; Palermos 2017). Another point of contact stems from the emphasis that AmI research places on the monitoring and anticipation of human behavior. This highlights the relevance of technologies that can be used to predict human behavior. The modeling and prediction of human behavior is, of course, a topic of long-standing interest within AI and cognitive science (Pentland and Liu 1999). Recently, however, such efforts have gained a new significance with the advent of predictive processing models of cognition (Clark 2016) and the extension of such principles to the design of predictive technologies (see Pezzulo et al. 2016).

3 The source of the smarts

As mentioned above, the IoT extends the conventional Internet to the realm of everyday objects and devices. Interestingly, this provides the basis for ever-more intimate forms of informational contact with the human social environment. In particular, the IoT provides additional opportunities (beyond those provided by the conventional Internet and Web) to observe and interact with human agents. Such forms of contact arguably yield an opportunity to extend the reach of machine intelligence (Smart 2017a, b), perhaps establishing the basis for socially-embedded (see Dautenhahn et al. 2002) or socially-situated (see Lindblom and Ziemke 2003) machines, i.e., machines that participate in complex forms of ‘perceptuomotor’ engagement with the (ambient³) human social environment.

The upshot of all this is a view of social machines that appeals to the role of the human social environment in shaping, supporting, and scaffolding the emergence of machine-based ‘cognitive’ capabilities. This is what we will call the *socially-situated machines view*:

Socially-Situated Machines View

Social machines are intelligent systems whose processing capabilities are oriented to the human social domain. Such systems are able to derive cognitive benefit from their contact with the human social environment.

What does it mean to say that a social machine is a system that derives cognitive benefit from the human social environment? And what, precisely, is the cognitive scientific relevance of this general idea? The answer (or at any rate, one answer) comes in the form of socially-scaffolded development and social learning.⁴

Turning first to the issue of scaffolded development, it has been suggested that human agents can contribute to the developmental emergence of machine-based cog-

³The use of the term “ambient” encourages us to see the Internet as affording access to the human social environment, thereby enriching the opportunities for machine-based cognizing (see Smart 2018). Thus, just as the Internet can be seen to contribute to a progressive enrichment of the technological and social environment in which human cognition occurs, so it can also be seen to enrich the environment in which some of our most advanced intelligent systems are now operating.

⁴Of course, the notions of socially-scaffolded development and social learning are not entirely distinct. Social learning, for example, might be one of the (many) processes that supports socially-scaffolded development.

nitive capabilities by participating in a form of ‘artificial cognitive ontogenesis’ (Dautenhahn 2007; Lindblom and Ziemke 2003). This idea has, of course, been one of the focal points of research into what is known as developmental robotics (see Dautenhahn and Billard 1999). But, we might ask, is there any reason why the same idea could not be applied to other kinds of intelligent system, such as the (non-robotic) systems that inhabit the large-scale networked niches of the Web, Internet, and the IoT? We suggest not. Indeed, we suggest that the IoT permits ever-more intimate forms of contact with the human social environment, providing insights into the dynamics of human behavior and creating new opportunities for cognitively-potent forms of human–machine interaction. A social machine, in this sense, can perhaps be thought of as the online equivalent of a developmental robotics system. It is, in other words, a system that is able to benefit, in a cognitive sense, from its immersion in an ontogenetic/developmental niche that is founded on Internet-enabled forms of contact with a significant chunk of humanity.⁵

In some ways, this idea is not particularly novel. Claims about the importance of social contact for machine intelligence date back to the early days of AI. Turing (1950), himself, raised the possibility of human care-givers playing a productive role in the emergence of advanced forms of machine intelligence. But the advent of the IoT does add a novel twist to this familiar idea. For the kind of social contact we envisage (i.e., the sort of contact provided by the IoT) is radically unlike anything that Turing (and many of his successors) could have anticipated. It is a form of contact in which the *global human community* is seen to play a role in supporting, shaping, and scaffolding the emergence of machine-based cognitive capabilities. This is important, for no one, we suspect, will want to be burdened with the task of nurturing a ‘child machine’ (see Turing 1950) to full cognitive maturity. Fortunately, the global reach of the Internet provides us with a potential superabundance of cognitive resources—a “cognitive surplus” (Shirky 2010)—which allows us to entertain the possibility that ‘parental’ responsibilities could be effectively ‘crowdsourced’ to large numbers of human individuals.

Next, let us consider the cognitive significance of the IoT from the perspective of social learning. In this case, the IoT might be seen to yield an opportunity to observe the human social environment and acquire information about various forms of human competence. The basic idea is that the Web, Internet, and IoT enable us to treat the human social environment as a source of information and knowledge, one that can be

⁵This is an idea that dovetails with work that is spread across a multiplicity of academic disciplines, including evolutionary psychology, cognitive science, and the philosophy of mind. In particular, the idea is relevant to work in developmental robotics (Dautenhahn 2007), Vygotskian views regarding the ontogenetic bases of human/machine intelligence (Lindblom and Ziemke 2003), and work regarding the centrality of the human social environment to the evolutionary emergence of the modern human mind (Flinn and Alexander 2007; Flinn et al. 2005; Sterelny 2007). If, for example, the human mind is something of a social artifact—a cognitive ‘device’ whose functional profile was shaped by the need to deal with the demands of a complex social environment—then perhaps Internet-enabled forms of informational contact with the human social world are poised to play a crucial (perhaps indispensable!) role in yielding AI systems whose cognitive capabilities echo those of their human counterparts (for more on this, see Smart and Madaan 2017).

mined and monitored as part of an overall effort to extend the cognitive and epistemic reach of certain forms of machine intelligence.

To help us understand this claim in a little more detail, consider the effort to develop self-driving cars. Such efforts clearly depend on advances in our ability to engineer sophisticated forms of sensor processing, especially in the visual domain. However, they also rely on advanced control systems that are able to respond in an intelligent manner to a multitude of road-relevant situations. In order to emulate the behavior of human road users, it thus seems important to capture at least some of the (commonsense) knowledge that human drivers have acquired as a result of their experience behind the wheel.⁶

How do we go about building cars that possess the behavioral competence and road-related *savoir faire* exhibited by the typical human driver? One possibility is to track the behavior of human-driven vehicles as they move around the road network, and then attempt to extract and formalize interesting regularities from the resultant body of ‘behavioral’ data. Such data sets are likely to be particularly valuable in cases where it is possible to track the precise behavior of vehicles at particular locations, such as at an intersection, a roundabout, or a notorious black spot. Additional value comes from the ability to track other kinds of information, such as the use of driver signaling mechanisms (e.g., the use of indicators and headlights) and information about prevailing meteorological conditions (e.g., the presence of fog).

The main point of this example is that it helps us see how a particular form of access to the human social environment can provide insight into bodies of experientially-grounded knowledge, some of which may be relevant to the attempt to engineer intelligent systems. The vision is thus one in which advanced forms of machine intelligence begin to emerge as the result of a novel opportunity to observe and learn about the human social environment. According to this vision, machine intelligence is, in some sense, parasitic on human experience: human experience supports the acquisition of particular forms of cognitive and behavioral competence that might be difficult to acquire via other means.

The socially-situated view of social machines thus helps us to locate the source of the smarts when it comes to the development of ever-more sophisticated forms of machine intelligence. Rather than see the ‘cognitive’ capabilities of such systems as stemming *solely* from the ingenuity of human engineers, the socially-situated view encourages us to adopt something of an *ecological* approach to machine intelligence (see Smart 2017b): it encourages us to see machine-based cognitive capabilities as arising from Internet-enabled forms of informational contact with the human social world.

4 Social observatories

An oft-mentioned virtue of the Internet is its capacity to improve our understanding of the social realm. Contemporary human societies, it should be clear, are massively

⁶Another issue relates to the codification of social norms pertaining to the tacitly accepted ‘rules of the road’ (see Walport 2014).

complex systems, ones whose scale and sophistication threatens to overwhelm our capacity for prediction, explanation, and scientific understanding. The emergence of fields such as computational social science help us see how such unruly and elusive beasts might be sufficiently tamed (or at least yoked) in such a way as to support their scientific analysis (Strohmaier and Wagner 2014). The basic idea, here, is that the advent of the Internet and Web, in conjunction with the deployment of IoT devices (especially sensors), provides opportunities to monitor social processes, yielding voluminous bodies of digital data that can be used to inform the scientific study of the human social environment.

Perhaps the ultimate expression of this idea comes in the form of the social observatory concept (Caton et al. 2015; Aharony et al. 2011), where a social observatory is a system that supports the observational analysis of society. This serves as the basis for a view of social machines that we will dub the *social observatory view*:

Social Observatory View

Social machines are systems that are intended to support the observational analysis of society at a variety of social scales. They are used to inform our understanding of social phenomena, typically for descriptive, predictive, or explanatory purposes.

As should be clear from this characterization, social machines (*qua* social observatories) are of crucial relevance to social science. According to the social observatory view, social machines are intended to improve our understanding of social phenomena, and this is something that pretty much defines the empirical and theoretical remit of the social sciences. This does not mean, however, that the social observatory view is irrelevant to the sciences of the mind. As was discussed in earlier sections (e.g., Section 2), cognitive systems can sometimes assume the form of multi-agent systems, where cognitive processing routines are distributed across a collection of cognitive agents. In such cases, it seems plausible to conclude that some cognitive phenomena are also social phenomena, at least insofar as we see social phenomena as tied to the interactions and exchanges that occur between the members of a multi-agent community. In this sense, social machines could assist with the effort to monitor the information processing dynamics of distributed or collective cognitive systems, expanding the palette of digital cognitive ethnographic techniques that are used in more restricted social settings (see Hulchins 2014).

A further link to the sciences of the mind comes from the attempt to apply machine learning techniques to bodies of social data. Of particular interest is the idea that computational architectures built around the principles of predictive coding or predictive processing might give rise to socially-oriented generative models that embody the causal structure of the human social environment (see Smart and Madaan 2017). Although speculative, this idea dovetails with the burgeoning of philosophical, scientific, and engineering research into deep learning techniques, predictive processing architectures, and generative models (see Clark 2016). It is, moreover, an idea that relates to the goals and interests of the social scientific community; for inasmuch as we see a generative model as capturing “the statistical structure of some set of observed inputs by inferring a causal matrix able to give rise to that very structure” (see Clark 2016, p. 41), then a socially-oriented generative model may provide insight into the forces and factors that shape the structure of the human social world.

Relative to this vision, the scientific significance of social machines (*qua* social observatories) is clear: such systems are the equivalent of biological sensors—just like biological sensors they provide the raw sensory data over which a hierarchically-organized economy of progressively more abstract cognitive(?) representations are constructed. The result may be a form of ‘deep’ understanding of human social behavior that parallels the sort of understanding that is deemed to arise in the case of brain-based forms of predictive processing (see Clark 2013).

5 The nuts and bolts of society

From the standpoint of the social observatory view, social machines are seen to play a crucial role in helping us improve our understanding of social phenomena. In fact, such understanding typically comes in a particular form, namely, an understanding of the mechanistic bases of social phenomena. In describing a particular form of technologically-mediated social observation, for example, Aharony et al. (2011) appeal to the idea that observation assists with the effort to identify the *mechanisms* that underlie social phenomena:

In the same way that fMRI techniques help map the interconnections and mechanics of the human brain, we hope that our work will help advocate an evolution from mostly passive observatories to data-rich Social fMRI type of studies that can help further our understanding of the interconnections and mechanics of human society. (Aharony et al. 2011, p. 658)

The search for mechanisms that underlie (or realize) phenomena is, of course, a goal that is shared by many scientific disciplines (Craver and Tabery 2016). When it comes to social phenomena, however, some have questioned whether an exclusive focus on the social realm (i.e., the realm of human agents) can carry the explanatory weight that upholds a scientific, mechanistically-informed understanding of society. Of particular concern is the extent to which the mechanisms that realize social phenomena are confined solely to the social realm. If this turns out *not* to be the case—if, for example, the relevant mechanisms involve material artifacts and technological devices in addition to social (i.e., human) agents—then an exclusive focus on the social realm would be to the overall detriment of the social sciences, undermining the extent to which a mechanistically-oriented explanatory account could be developed for perhaps even the most straightforward of social phenomena (Latour 1992; Orlikowski and Scott 2008). The result is that social phenomena would tend to resist explanation in terms of purely social mechanisms and thus look somewhat ‘mysterious’. In this respect, Latour (1992) suggests that trying to explain the social in the absence of the technological is akin to focusing attention on one half of a tennis court during a tennis match: in the absence of an appreciation of what is happening on the other side of the net the observed phenomena appear “as so many meaningless moves” (p. 168).

A similar sort of worry lies at the heart of recent appeals to active externalism within the philosophy of mind (Clark and Chalmers 1998; Clark 2008). Active externalism is the idea that the mechanistic supervenience base for human mental states

and processes extends to include resources that lie external to the biological borders of skin and skull. The implications of this idea for cognitive science are pretty much the same as those identified (by Latour) for social science. The point is that if you want to develop a scientific (mechanistically-grounded) understanding of some phenomenon of interest (e.g., a cognitive or social phenomenon), then it pays to consider the possibility that the mechanisms responsible for the phenomenon in question could involve a rather surprising set of material elements, ones that perhaps lie beyond the traditional analytic orbit of a given scientific discipline. In seeking to explain some phenomenon, therefore, we should be wary of a form of infelicitous boundary placement concerning the mechanistic limits of a system. As was noted by the cyberneticist Bateson (1972, p. 459), in delineating a system, it is imperative that we do not impose borders and boundaries that cut across the links and connections of mechanistically-relevant circuits. To do so, Bateson warns, risks leaving us with a set of phenomena that may very well lie beyond our explanatory grasp.

By heeding such warnings, we arrive at a view of social machines that appeals to issues of mechanistic explanation and mechanistic realization (see Wilson and Craver 2007). This is the *mechanistic view* of social machines:

Mechanistic View

Social machines are systems whose events, states, and processes are realized by socio-technical mechanisms. In other words, social and technological elements form part of the physical machinery that realizes system-level phenomena.

In essence, the mechanistic view is the idea that a social machine is home to a hybrid mechanism that consists of both social (i.e., multiple human individuals) and technological (e.g., IoT devices) components. The social and technological components must, in this case, be jointly relevant to the realization of phenomena in which the larger systemic (socio-technical) organization is involved.

One advantage of the mechanistic view is its compatibility with (at least some) approaches in the social sciences. The mechanistic view, for example, speaks to the interests and concerns of the proponents of analytical sociology (Hedström 2005; Hedström and Ylikoski 2010), and it establishes sensible contact with approaches that draw attention to the entanglement of social and technological (or at least material) elements in the course of everyday social life (see Orlikowski and Scott 2008).

The mechanistic view is also consistent with much of the work in the sciences of the mind, where the aim, in many cases, is to develop a better understanding of the *mechanisms* that underlie cognitive phenomena (Bechtel 2008; Craver 2007b). It is, in particular, a view that is perfectly compatible with what we will call the *cognitive systems view* of social machines—the idea that social machines are a form of distributed or collective cognitive organization:

Cognitive Systems View

A social machine is a form of distributed cognitive system in which cognitive processing routines are realized as a result of the interactions between (multiple) human individuals and a collection of technological (Web- or Internet-based) components.

The cognitive systems view has recently been championed by Palermos (2017). Palermos suggests that we should see social machines as cognitive systems that comprise both social and technological elements. In other words, social machines (*qua* cognitive systems) are systems in which system-level cognitive phenomena are realized by the interactions and exchanges that occur between the elements of an underlying socio-technical fabric. The cognitive systems view and the mechanistic view thus share a common commitment to the idea that social machines are hybrid (socio-technical) organizations, and underlying this hybridity is a commitment to issues of constitutive relevance (see Craver 2007a): the fact that *both* social and technological components are involved in (are constitutively relevant to) the mechanistic realization of certain kinds of phenomena. It is, in fact, this commitment to constitutive relevance that determines the hybrid status of a social machine. For if the phenomenon of interest were to be realized solely by (say) the technological components, then the social components (i.e., the human agents) would not be part of the mechanism that realized the focal phenomenon. The result would be that our mechanistic understanding of the relevant phenomenon would not need to appeal to the social domain, and we could thus dispense with claims about the hybrid nature of the relevant system and its associated mechanisms.

Despite the compatibility between the mechanistic and cognitive systems views, there is an important difference between them. This difference concerns the nature of the phenomena that are realized by the interactions between a collection of social and technological elements. The mechanistic view is utterly neutral on this point: it makes no claims about the nature of the phenomena that are exhibited by a socio-technical system (and thus realized by materially hybrid socio-technical mechanisms). The cognitive systems view, in contrast, is much more specific. It insists that social machines are to be individuated with respect to the nature of the phenomena (i.e., cognitive phenomena) that are exhibited by the target system. According to this characterization, the cognitive systems view may be seen as a more restricted version of the mechanistic view (or, perhaps better, a specialization of the mechanistic view). In other words, one way of understanding the relationship between the cognitive systems view and the mechanistic view is to see the social machines picked out by the cognitive systems view as corresponding to a proper subset of those picked out by the mechanistic view.

Is there any reason to prefer the mechanistic view over the cognitive systems view? Given that the mechanistic view is a broader concept—one that encompasses the cognitive systems view—it should be clear that nothing is lost if we opt to embrace the mechanistic view. It is less clear that the same is true if we opt for the cognitive systems view. In this case, our focus of interest is restricted to a class of systems that are specifically associated with phenomena of the cognitive kind. Setting aside the not insignificant problem of determining when a phenomenon should be counted as a *bona fide* cognitive phenomenon (Adams and Aizawa 2010; Clark 2010), it is not entirely clear why the science of social machines should restrict its focus of interest in this way. What is it, we might ask, that is so special about cognitive phenomena that they should serve as the basis of a new multi-disciplinary endeavor, especially

when the empirical and theoretical targets of that endeavor are already the subject of an existing and well-established area of cognitive scientific research (i.e., distributed cognitive science)? In addition to this, we might ask why it is that the impact of the Internet, and its extension into the physical world via the IoT, should be limited solely to the realm of cognitive phenomena. Are there not some types of phenomena—such as those associated with the domains of economics, sociology, and political science—that might be of equal (if not exclusive!) interest to the social sciences rather than just the cognitive sciences?⁷ Proponents of the cognitive systems view may, of course, be inclined to argue that all these phenomena (irrespective of their economic, sociological, or political nature) are ultimately grounded in a set of cognitive mechanisms (i.e., mechanistic explanations for social phenomena will, at some point or another, have to appeal to mechanisms that realize cognitive phenomena). Philosophically, however, this a dangerous move. For the issue at stake is not whether some form of cognitively-relevant machinery is merely a constituent of a larger system. Rather, the question is whether the mechanisms that realize the phenomena of interest are of the hybrid variety, in the sense of straddling *both* the technological and social domains. If this should turn out not to be the case—if, for example, all the cognitively-relevant processing within a socio-technical system was to be realized by mechanisms located within the heads of the human constituents of the relevant system—then we would not have a genuine form of distributed cognition, and thus we would not have, according to the cognitive systems view, a *bona fide* social machine.

None of this should be seen as undermining the legitimacy of the cognitive systems view, for there are no doubt times when it *does* make sense to view a social machine as a form of distributed cognitive organization. In these situations, it seems perfectly appropriate to view the relevant system through the conceptual lens provided by the cognitive systems view. The key question, therefore, is not the validity of the cognitive systems view; rather, it is whether we should allow the theoretical and empirical remit of the science of social machines to be limited to just those systems that are involved in phenomena of the cognitive kind. Our own view is that the answer to this question is ‘no’, and in the absence of a reason to reject the mechanistic view, it is far from clear that there is any need for further discussion on the matter. The mechanistic view happily accommodates the idea that some (and perhaps many) social machines will be usefully understood (and fruitfully explored) as cognitive systems. At the same time, it allows for the possibility that social machines may, on occasion, be involved in phenomena that are not so easily characterized as cognitive in nature. This looks to be important when it comes to the multi-disciplinary nature of the social machine research effort. In particular, by directing attention to the nature of the mechanisms that realize system-level phenomena, the mechanistic view establishes a philosophically-important link to mechanistic approaches in the social (Ylikoski

⁷Examples might include cryptocurrency price fluctuations, the stability/instability of a political regime, a protest movement, and so on.

2018), cognitive (Zednik 2018), and computational (Piccinini 2015b) sciences, and it does so in such a way as to fully accommodate the idea that social machines may, at times, be studied, analyzed, and conceptualized as (distributed) cognitive systems.⁸

6 Socially-extended computation

The IoT promises to deliver a world that is littered with a rich array of networked computational artifacts, and this enables us to think of the physical environment as, in some sense, a computational resource that is subject to programmatic control (see Taivalsaari and Mikkonen 2017). Such ideas lie at the heart of a number of research efforts in contemporary computer science (especially research into ubiquitous, pervasive, and mobile computing). They also serve as the basis for a number of visionary concepts and research challenges, such as the notion of the *global ubiquitous computer* (see Kwiatkowska et al. 2004).

But it is not just the growing proliferation of computational devices that has piqued the interests of the computer science community. Increasingly, computer scientists are concerned with the extent to which the human social environment can, itself, be treated as a computational resource. Although research in this area goes by a variety of names—including social computation (Kearns 2012), human computation (Michelucci 2013), the social computer (Robertson and Giunchiglia 1987), and the crowd computer (Kucherbaev et al. 2012)—a common focus of attention is the extent to which human agents can be factored (or incorporated) into computational routines, thereby yielding computational organizations in which human agents are considered “as part of the hardware of the system and [are] available for the realization of computational tasks” (Jara et al. 2013, p. 10).

All of this, we suggest, provides us with another way of thinking about social machines. In this case, social machines are cast as hybrid computational systems that

⁸It is worth emphasizing this particular point, since the nature of the relationship between the mechanistic view and the cognitive systems view was the source of much controversy during the review stage of the present paper. Our own view, here, is that it is a mistake to see the mechanistic view and the cognitive systems view as competing positions when it comes to our understanding of social machines. The mechanistic view, we suggest, is best seen as a view that subsumes the cognitive systems view (and other views) by virtue of the fact that it does not tie itself to claims about the specific nature of the phenomena that are realized by socio-technical mechanisms. This allows for the possibility that social machines will be involved in a variety of disparate phenomena (e.g., social, computational, cognitive, economic, and political phenomena), while avoiding a number of philosophical puzzles and problems (e.g., those relating to the ‘mark of the cognitive’ (Adams 2010; Adams and Garrison 2013)). The mechanistic view and the cognitive systems view are only in competition if proponents of the cognitive systems view insist that *all* social machines are to be treated as cognitive systems. To our mind, this seems overly restrictive, and it risks losing contact with what is a genuinely multi-disciplinary effort surrounding a rich array of online systems. To be sure, the cognitive systems view plays an important role in alerting us to particular features and issues when we are dealing with certain kinds of social machine (i.e., social machines of the cognitive kind), but other specializations of the mechanistic view (e.g., the socio-computational view discussed in Section 6) will also have their own specific virtues. The point of the mechanistic view, here, is that it provides us with a common theoretical framework for understanding the relative merits of particular views in particular situations.

include multiple human individuals as part of the physical machinery that realizes a computational process. This is what we will call the *socio-computational view* of social machines:

Socio-Computational View

Social machines are hybrid computational systems, in the sense that they consist of both social and technological elements. These elements are jointly involved in the realization of computational processes (or other computational phenomena).

At first sight, this particular view of social machines seems to have a lot going for it. One advantage of the socio-computational view is that it helps to establish a point of contact with recent work in computer science, especially that relating to social computation, human computation, crowd computing, and so on. In casting social machines as a particular form of computational system, the socio-computational view thus looks to be ideally placed to speak to the interests of the computer science community.

The socio-computational view is also a view that speaks to the notion of *heterotic computing* (Kendon et al. 2015; Horsman 2015). In this case, computational benefits are seen to derive from the amalgamation of two distinct, albeit complementary, computational substrates. A heterotic computing system is thus a materially hybrid organization that seeks to exploit the distinctive computational and representational properties of its constituent elements. Some notable examples include neurocomputers (Ruaro et al. 2005), Physarum machines (Adamatzky 2010), and, according to the socio-computational view, social machines (Horsman 2015; Horsman et al. 2017).

Finally, the socio-computational view touches on a number of topics that have dominated recent discourses within the philosophy of mind and cognitive science. These include the aforementioned notion of active externalism (see Section 5), which is the idea that human cognitive processes can, on occasion, be subject to non-standard forms of mechanistic realization involving both the biological brain and a motley assortment of extra-organismic elements (Clark and Chalmers 1998; Clark 2008). The socio-computational view of social machines builds on this basic idea, helping us see social machines as something akin to socially-extended computational systems, i.e., systems in which human agents form part of the (non-standard) realization base for conventional(?) forms of computational processing.

Some of the problems afflicting the socio-computational view resemble those discussed in relation to the cognitive systems view (see Section 5). As with the cognitive systems view, the socio-computational view encourages us to think of social machines as a particular kind of socio-technical system, namely, one that implements a computational process, or (perhaps more generally) realizes a computational phenomenon. The problem, however, is that our understanding of what it is that makes a process computational (as opposed to non-computational) is surprisingly poor (see Piccinini 2015a). Neither is it clear why we should restrict our attention solely to the realms of the computational. In limiting our attention to the domain of purely computational systems, therefore, we encounter the by now familiar (see Section 5) problems of ambiguity (how do we individuate computational/cognitive phenomena from non-computational/non-cognitive phenomena?)

and restricted scope (why should we limit our attention solely to the realm of computational/cognitive phenomena?).⁹

The mechanistic view discussed in Section 5 provides a potential solution to these problems. The mechanistic view, recall, sidesteps the problem of how to individuate specific phenomena (e.g., how to limn the realm of the cognitive or the computational) by focusing attention on the nature of the mechanisms that realize system-level phenomena. From the standpoint of the mechanistic view, what matters is the *way* in which phenomena are realized, not the nature of the realized phenomena. In particular, the claim is that system-level processes are realized by socio-technical mechanisms, particularly those that involve the use of Internet communication protocols. In this sense, the mechanistic view is poised to accommodate the socio-computational view in the same way it accommodates the cognitive systems view: it allows for the possibility that social machines may be studied as computational (or cognitive systems) without thereby becoming embroiled in debates about how computational (or cognitive) phenomena should, themselves, be understood.

7 Social machinery

All of the views discussed thus far have focused on either the properties of the technological components of a socio-technical system (e.g., the socially-situated view) or the properties of the socio-technical system itself (e.g., the mechanistic view). None of the views, however, have focused attention on the social elements of a socio-technical system. This looks to be a significant shortcoming, for there is surely a sense in which both current and future technologies are poised to influence the dynamics of human social behavior. What impact, if any, do such influences have on our attempt to understand and study social machines?

One possibility is that the advent of technologies like the Web, the Internet, and the IoT will alter the way that society is itself viewed. We might thus be inclined to view society in increasingly ‘machine-like’ terms by virtue of an enhanced capacity to observe and analyze social processes (see Section 4). The inspiration for this view comes from work relating to the *quantified self* (Swan 2013). Lupton (2013), for example, claims that the widespread availability of portable, network-enabled devices (e.g., wearable devices) transforms our view of the human biological body. The quantified-self movement, Lupton (2013) suggests, “begins to construct a view of the body/self as a machine-like entity, with ‘inputs’ and ‘outputs’ . . . that can readily be measured and quantified. . . . The body in this discourse becomes positioned as a ‘smart machine’ interlinked to other ‘smart machines’ . . .” (pp. 26–27). By extending this idea to the social realm, we encounter the idea of society as something like a machine by virtue of the quantification opportunities afforded by the IoT.

⁹Additional issues are no doubt likely to arise in respect of the relationship between social machines of the cognitive kind and social machines of the computational kind: Are these the same systems, for example? Are all computational phenomena *bona fide* cognitive phenomena? And so on.

An alternative way in which society may come to acquire a machine-like character is via the technological shaping of social processes. Consider, for example, the way in which a broad array of social processes now involve the use of digital technologies. We rely on these technologies to such an extent that it is perhaps difficult to imagine certain kinds of social processes being undertaken in the absence of advanced forms of computational control and coordination. Accordingly, we may start to conceive of society as something like a machine by virtue of what might be called the ‘mechanization’ of social processes, i.e., the use of computational technology to control, organize, and manage the social realm.

What we end up with, therefore, is a view of social machines that reflects the impact of technological systems (including IoT devices) on the properties of social systems. This is a view that we will dub the *society as a machine view*:

Society as a Machine View

Social machines are social systems (i.e., systems consisting of multiple social agents) that exhibit properties consistent with those of a machine. These machine-like properties arise as the result of the interplay between social and technological elements.

One of the immediate problems faced by this view is the ambiguous nature of the term “machine.” It might be claimed, for example, that society is radically unlike the kinds of systems that we typically associate with the term “machine” (e.g., a Jacquard loom or a digital computer). And this perhaps remains the case, irrespective of whatever sorts of instrumentation and measurement are brought to bear on the social realm.

This is precisely the sort of criticism that is leveled at the social machine concept by McBride (2011). He suggests that:

The metaphor of a machine carries with it ideas of managerial control where through rewards and penalties control of a group of people . . . can be regulated. It carries with it also the idea of regularity and repeatability. In pursuing the idea of a social machine we are suggesting that interactions among the people and the technology are predictable and controllable. (McBride 2011, p. 2)

In this respect, McBride (2011) finds the social machine concept wanting. The inherent unpredictability and unruly nature of the social realm prohibits, he suggests, the development of computational specifications that succeed in capturing the complexity of human–technology interactions. This undermines the legitimacy of machine-like metaphors, and, in the worst case, threatens the integrity of the social machine concept itself.

The basic problem, then, is that McBride associates the concept of a machine with the possession of certain properties (namely, predictability and controllability). Human social systems, he suggests, lack these properties, and the result is a rejection of the society as a machine view. The structure of this argument is as follows:

- **Premise 1:** We individuate machines with respect to a set of behavior-related properties. In particular, a machine is a system that exhibits the properties of predictability and controllability.

- **Premise 2:** The behavior of human social systems is neither predictable nor controllable.
- **Conclusion:** Therefore, social systems are not machines, and machine-based metaphors are inappropriate ways of characterizing human social behavior and/or human–technology interactions.

One response to all this is to question the relevance of predictability and controllability in determining what it is that makes something a machine (i.e., to challenge Premise 1 above). Here, it will help to revisit a well-known parable, first presented by one of the towering figures of AI: Herbert Simon.

The Parable of the Ant

We watch an ant make his laborious way across a wind- and wave-molded beach. He moves ahead, angles to the right to ease his climb up a steep dune let, detours around a pebble, stops for a moment to exchange information with a compatriot. Thus he makes his weaving, halting way back to his home. . . Viewed as a geometric figure, the ant’s path is irregular, complex, hard to describe. But its complexity is really a complexity in the surface of the beach, not a complexity in the ant. (Simon 1996, p. 51)

The parable of the ant is important because it highlights the crucial role of the *environment* in determining the complexity of even (perhaps) the simplest of systems. Consider, for example, the kinds of systems described by Braitenberg (1984) in his wonderfully engaging and insightful book, *Vehicles: Experiments in Synthetic Psychology*. Braitenberg asks us to imagine a set of relatively simple robotic systems (i.e., machines). He then asks us to anticipate the behavior of such systems once they are embedded in a sufficiently rich and well-structured environment. The result, Braitenberg (1984) suggests, is a pattern of behavior whose complexity owes as much to the features of the external environment as it does to the properties of the internal control mechanism.

Such insights, we suggest, are poised to do substantial damage to claims about the importance of controllability and predictability in determining when something should be counted as a machine. As empirical explorations of Braitenberg’s work show (see Dawson et al. 2010), the behavior of a simple mechanical system (e.g., a LEGO robot equipped with a computational control system) can, on occasion, exhibit surprisingly complex (and unpredictable) behavior, and this complexity often varies with respect to features of the external environment. Does this mean that a robot is a machine in one situation, but not another? Does the robot’s status as a machine come and go as it moves from one part of the environment to the next? This looks to be unlikely. And the result, we suggest, is that predictability and controllability are unlikely to be of crucial importance in determining what it is that makes something a machine.

The best way to understand the notion of a machine, we suggest, is to focus our attention on issues of mechanistic realization (see Wilson and Craver 2007). In particular, we suggest that the term “machine” denotes a system that contains mechanisms that realize phenomena that are associated with the system in question. Inasmuch as we accept this to be the case, then a social machine is perhaps best seen as a collection of mechanisms that realize system-level phenomena. The real issue, then, relates

to how we individuate a social machine from other (non-social) machines: what is it about a *social* machine that distinguishes it from machines of the non-social variety? The answer to this question may lie in either the nature of the mechanisms themselves or the kind of phenomena that are realized by the mechanisms. A social machine might thus be a system (any system!) that contains social mechanisms (e.g., mechanisms consisting of multiple social agents). On the other hand, a social machine might be regarded as a system (any system!) that realizes social phenomena. Perhaps, in the final analysis, the choice between these alternatives is merely a false dilemma. It is by no means clear, for instance, that social phenomena can arise in the absence of social mechanisms. Neither is it clear that the phenomena realized by a social mechanism can be anything other than social phenomena.

8 Discussion: a hot supper?

8.1 The Internet of Things

One of the aims of the present paper was to assess how the IoT informs philosophical debates associated with the science of social machines. The majority of work relating to social machines has tended to limit its attention to the realm of Web-based systems, and it is therefore easy to think that the scientific targets of the social machine research effort should be limited to the study of social networking systems (e.g., Facebook), citizen science platforms (e.g., Galaxy Zoo), collaborative authoring environments (e.g., Wikipedia), and so on. The results of the present work lead us to question this assumption. We have shown, for example, how a consideration of

Table 1 The impact and relevance of the IoT for the science of social machines. The cognitive systems view and the socio-computational view inherit the properties of the mechanistic view and are not included in the table

View	IoT impact
Ambient intelligence view	IoT devices provide new opportunities to support the activity of multiple individuals from an AmI perspective.
Socially-situated machines view	IoT devices support unprecedented levels of contact with the human social environment. In doing so, they provide new opportunities to learn about human behavior and exploit human knowledge.
Social observatory view	IoT devices provide new opportunities to observe and study human behavior at a variety of social scales.
Mechanistic view	IoT devices serve as part of the material fabric for novel socio-technically hybrid mechanisms. In effect, IoT devices become incorporated into the mechanisms that realize perhaps entirely new kinds of cognitive, social, and computational phenomena.
Society as a machine view	IoT devices provide new opportunities to influence human social behavior and thus change society.

the IoT proffers both a philosophical and scientific payoff—broadening the potential scope of the social machine research effort and providing us with new ways to think about the meaning of the term “social machine.” The views presented in this paper are testament to this sort of payoff. In some cases, the IoT provides us with new opportunities (and challenges) for research in specific areas, while in other cases a particular view of social machines is *only* possible if we survey the conceptual terrain from the specific vantage point provided by the IoT (see Table 1). (The ambient intelligence view, for example, is only possible if we accept that the IoT should form part of our philosophical thinking about social machines.)

None of this should lead us to question the validity of existing work on social machines, specifically that which limits its attention to the realm of Web-based systems. It is even possible that proponents of a purely Web-based view of social machines could reap the rewards of the present analysis without forsaking a selective focus on the Web. There is, for example, a Web-based counterpart to the IoT—the Web of Things (WoT) (Guinard and Trifa 2009; Guinard et al. 2011)—that is perfectly able to absorb the benefits of the present analysis. The key question to ask in this case, however, is why the science of social machines should be restricted to a specific set of Internet communication protocols: What do we lose if we broaden the scope to include systems built on top of Internet protocols other than those associated with the Web?

8.2 The sciences of the mind

Much (although by no means all) of the social machine research effort is concerned with the social impact of social machines, e.g., their implications for social change and effects on social policy. This might lead us to conclude that the sciences of the mind (i.e., cognitive science, AI, philosophy of mind, neuroscience, and so on) are of little relevance to the science of social machines. The present paper challenges this conclusion. The ambient intelligence view (see Section 2) and the cognitive systems view (see Section 5), for example, help to reveal the links between social machines and research into distributed cognitive systems. Other points of interdisciplinary contact concern the social shaping of machine-based cognitive abilities (see Sections 3 and 4), the use of predictive processing techniques to yield a ‘deep’ understanding of human social processes (see Section 4), and the extension of active externalist theorizing to the realm of socio-computational systems (see Section 6).

8.3 The definitional endeavor

The present analysis reveals a number of different views of social machines (see Table 1), and these views sit alongside those that have already been discussed within the philosophical and scientific literature, e.g., the content creation view (Berners-Lee and Fischetti 1999), the cognitive systems view (Palermos 2017), and the problem-solving view (Hooper et al. 2016). What is the impact of this efflorescence of different views for our conceptual understanding of social machines?

One option, perhaps, is to abandon the search for a canonical definition and accept that the term “social machine” is applicable to multiple systems. One benefit of this

sort of approach is that it enables us to use the term “social machine” as something akin to a conceptual instrument that helps to bring a range of different systems and different properties into sharper focus at different times. In this sense, different views of social machines would play a role similar to that fulfilled by the spectral filters used in (e.g.) astronomical research. In essence, we might be able to see the different views as affording a *multi-spectral conceptual analysis* of the emerging socio-technical landscape, with each view helping to highlight specific structural and functional features that are perhaps invisible, or at least difficult to resolve, from an alternative perspective.

A second option is to insist that the science of social machines is best served by a single view, one that helps to ensure scientific activity is undertaken in a coherent, coordinated, and theoretically-integrated manner. In this respect, the mechanistic view (see Section 5) may be of particular interest. Not only does this view subsume two other views (i.e., the cognitive systems view and the socio-computational view); it is also relevant to a number of other views discussed in the present paper. The ambient intelligence view (see Section 2), for example, could easily be cast as a variant of the cognitive systems view,¹⁰ and thus compatible with the mechanistic view. In addition, the discussion associated with the society as a machine view (see Section 7) highlighted a dilemma that appeared to hinge on the acceptability of the mechanistic view: either one embraces McBride’s (2011) critique of the social machine concept, in which case one arrives at the conclusion that the social machine concept is flawed; or one embraces an alternative (mechanistically-oriented) conception of what it means for something to be a machine, in which case we are (at the very minimum) in the rough conceptual vicinity of the mechanistic view.

The mechanistic view of social machines is thus a view that is broadly compatible with a range of other views that have surfaced in the scientific and philosophical literature. It is, moreover, a view that is poised to highlight the social and cognitive scientific relevance of social machines by drawing on a rich array of familiar philosophical concepts. These include the concepts of mechanistic realization (Wilson and Craver 2007), mechanistic explanation (Craver and Tabery 2016), cognitive integration (Palermos 2017), constitutive relevance (Craver 2007a), and active externalism (Clark 2008). By applying these concepts to the realm of social machines we are in a favorable position, we suggest, to integrate a number of (otherwise disparate) sociological (Latour 1992; Hedström 2005) and cognitive scientific (see Palermos 2017) ideas within a common (philosophically-informed) conceptual framework.

An additional merit of the mechanistic view is the way it guides our thinking about the importance of the social machine research effort. To help us see this, consider one of the merits of the society as a machine view: its ability to highlight the social impact of technology and thus the social responsibilities of technology designers. As noted by Tromp et al. (2011):

Designers [can] no longer can hide behind the needs and wishes of the consumer; instead, they have to take responsibility as ‘shapers’ of society. Doing so

¹⁰In the sense that an AmI system can be seen as a form of distributed cognitive organization.

entails a shift from a user-centered approach to a society-centered one. (Tromp et al. 2011, p. 19)

This is an important point, but it is one that is no less applicable to the mechanistic view as it is to the society as a machine view. Inasmuch as we see the ever-expanding realm of the Internet as part of the physical machinery of the social world—part of the realization base for (e.g.) social phenomena—then the development of new Internet technology is no less a form of ‘social engineering’ as it is technological engineering. In developing new Internet technologies, we extend the palette of physical elements that can be used as the building blocks (the mechanistic realizers) of our future society. Perhaps it is this image, more than any other, that best captures the value and importance of the emerging science of social machines.

Acknowledgments This work is supported under SOCIAM: The Theory and Practice of Social Machines. The SOCIAM Project is funded by the UK Engineering and Physical Sciences Research Council (EPSRC) under grant number EP/J017728/1 and comprises the Universities of Southampton, Oxford and Edinburgh. Additional support was provided by the UK EPSRC as part of the PETRAS IoT Research Hub—Cybersecurity of the Internet of Things under grant number EP/N022912/1.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Aarts, E.H.L., & Encarnaç o, J.L. (2008). *True visions: the emergence of ambient intelligence*. Berlin: Springer.
- Adamatzky, A. (2010). *Physarum machines: computers from slime mould*. London: World Scientific Publishing Co. pte Ltd.
- Adams, F. (2010). Why we still need a mark of the cognitive. *Cognitive Systems Research*, 11(4), 324–331.
- Adams, F., & Aizawa, K. (2010). Defending the bounds of cognition. In Menary, R. (Ed.) *The extended mind* (pp. 67–80). Cambridge: MIT Press.
- Adams, F., & Garrison, R. (2013). The mark of the cognitive. *Minds and Machines*, 23(3), 339–352.
- Aharony, N., Pan, W., Ip, C., Khayal, I., Pentland, A. (2011). Social fMRI: investigating and shaping social mechanisms in the real world. *Pervasive and Mobile Computing*, 7(6), 643–659.
- Atzori, L., Iera, A., Morabito, G. (2010). The Internet of Things: a survey. *Computer Networks*, 54(15), 2787–2805.
- Bateson, G. (1972). *Steps to an ecology of mind*. London: Jason Aronson Inc.
- Bechtel, W. (2008). *Mental mechanisms: philosophical perspectives on cognitive neuroscience*. New York: Lawrence Erlbaum Associates.
- Berners-Lee, T., & Fischetti, M. (1999). *Weaving the web: the original design and ultimate destiny of the world wide web*. New York: Harper Collins.
- Berners-Lee, T., Hall, W., Hendler, J., Shadbolt, N., Weitzner, D.J. (2006). Creating a science of the Web. *Science*, 313(5788), 769–771.
- Braitenberg, V. (1984). *Vehicles: experiments in synthetic psychology*. Cambridge: MIT Press.
- Caton, S., Hall, M., Weinhardt, C. (2015). How do politicians use Facebook? An applied Social Observatory. *Big Data & Society*, 2(2), 1–18.
- Clark, A. (2008). *Supersizing the mind: embodiment, action, and cognitive extension*. New York: Oxford University Press.

- Clark, A. (2010). Coupling, constitution and the cognitive kind: a reply to Adams and Aizawa. In Menary, R. (Ed.) *The extended mind* (pp. 81–99). Cambridge: MIT Press.
- Clark, A. (2013). Expecting the world: perception, prediction, and the origins of human knowledge. *The Journal of Philosophy*, 110(9), 469–496.
- Clark, A. (2016). *Surfing uncertainty: prediction, action and the embodied mind*. New York: Oxford University Press.
- Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, 58(1), 7–19.
- Cooke, N.J., Gorman, J.C., Winner, J.L. (2007). Team cognition. In Durso, F.T., Nickerson, R.S., Dumais, S.T., Lewandowsky, S., Perfect, T.J. (Eds.) *Handbook of applied cognition*, 2nd Edn. (pp. 239–268). Chichester: Wiley.
- Craver, C. (2007a). Constitutive explanatory relevance. *Journal of Philosophical Research*, 32, 3–20.
- Craver, C. (2007b). *Explaining the brain: mechanisms and the mosaic unity of neuroscience*. Oxford: Clarendon Press.
- Craver, C., & Tabery, J. (2016). Mechanisms in science. In Zalta, E.N. (Ed.) *The stanford encyclopedia of philosophy*, Spring 2016 Edn. Stanford: Stanford University.
- Dautenhahn, K. (2007). A paradigm shift in artificial intelligence: why social intelligence matters in the design and development of robots with human-like intelligence. In Lungarella, M., Iida, F., Bongard, J., Pfeifer, R. (Eds.) *50 Years of artificial intelligence: essays dedicated to the 50th anniversary of artificial intelligence* (pp. 288–302). Berlin: Springer.
- Dautenhahn, K., & Billard, A. (1999). Studying robot social cognition within a developmental psychology framework. In Schweitzer, G., Burgard, W., Nehmzow, U., Vestli, S.J. (Eds.) *Third European workshop on advanced mobile robots* (pp. 187–194). Zürich: IEEE.
- Dautenhahn, K., Ogden, B., Quick, T. (2002). From embodied to socially embedded agents—implications for interaction-aware robots. *Cognitive Systems Research*, 3(3), 397–428.
- Dawson, M.R., Dupuis, B., Wilson, M. (2010). *From bricks to brains: the embodied cognitive science of LEGO robots*. Edmonton: AU Press.
- Flinn, M.V., & Alexander, R. (2007). Runaway social selection in human evolution. In Gangestad, S.W., & Simpson, J.A. (Eds.) *The evolution of mind: fundamental questions and controversies* (pp. 249–255). New York: The Guilford Press.
- Flinn, M.V., Geary, D.C., Ward, C.V. (2005). Ecological dominance, social competition, and coalitionary arms races: why humans evolved extraordinary intelligence. *Evolution and Human Behavior*, 26(1), 10–46.
- Greengard, S. (2015). *The internet of things*. Cambridge: MIT Press.
- Guinard, D., & Trifa, V. (2009). Towards the Web of Things: Web mashups for embedded devices. In *2nd workshop on mashups, enterprise mashups and lightweight composition on the web*. Madrid.
- Guinard, D., Trifa, V., Mattern, F., Wilde, E. (2011). From the Internet of Things to the Web of Things: resource-oriented architecture and best practices. In Uckelmann, D., Harrison, M., Michahelles, F. (Eds.) *Architecting the internet of things* (pp. 97–129). Berlin: Springer.
- Hedström, P. (2005). *Dissecting the social: on the principles of analytical sociology*. Cambridge: Cambridge University Press.
- Hedström, P., & Ylikoski, P. (2010). Causal mechanisms in the social sciences. *Annual Review of Sociology*, 36, 49–67.
- Hendler, J., & Berners-Lee, T. (2010). From the Semantic Web to social machines: a research challenge for AI on the World Wide Web. *Artificial Intelligence*, 174, 156–161.
- Hendler, J., Shadbolt, N.R., Hall, W., Berners-Lee, T., Weitzner, D. (2008). Web science: an interdisciplinary approach to understanding the Web. *Communications of the ACM*, 51(7), 60–69.
- Heylighen, F. (2013). From human computation to the global brain: the self-organization of distributed intelligence. In Michelucci, P. (Ed.) *Handbook of human computation* (pp. 897–909). New York: Springer.
- Hollan, J., Hutchins, E., Kirsh, D. (2000). Distributed cognition: toward a new foundation for human-computer interaction research. *ACM Transactions on Computer-Human Interaction*, 7(2), 174–196.
- Hooper, C., Bailey, B., Glaser, H., Hendler, J. (2016). Social machines in practice: solutions, stakeholders and scopes. In Nejdil, W., Hall, W., Parigi, P., Staab, S. (Eds.) *8th international ACM web science conference* (pp. 156–160). Hannover: ACM.
- Horsman, D., Kendon, V., Stepney, S. (2017). The natural science of computing. *Communications of the ACM*, 60(8), 31–34.

- Horsman, D.C. (2015). Abstraction/representation theory for heterotic physical computing. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 373, 20140224.
- Hutchins, E. (1995). *Cognition in the wild*. Cambridge: MIT Press.
- Hutchins, E. (2014). The cultural ecosystem of human cognition. *Philosophical Psychology*, 27(1), 34–49.
- Jara, J., Daniel, F., Casati, F., Marchese, M. (2013). From a simple flow to social applications. In Sheng, Q.Z., & Kjeldskov, J. (Eds.) *13th international conference on web engineering: current trends in web engineering* (Vol. 8295, pp. 39–50). Springer, Aalborg, Lecture Notes in Computer Science.
- Kearns, M. (2012). Experiments in social computation. *Communications of the ACM*, 55(10), 56–67.
- Kendon, V., Sebald, A., Stepney, S. (2015). Heterotic computing: past, present and future. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 373, 20140225.
- Kennedy, J.B. (1926). When woman is boss: an interview with Nikola Tesla. <http://www.tfcbooks.com/tesla/1926-01-30.htm>. Accessed 24 April 2017.
- Kucherbaev, P., Tranquillini, S., Daniel, F., Casati, F., Marchese, M., Brambilla, M., Fraternali, P. (2012). Business processes for the crowd computer. In La Rosa, M., & Soffer, P. (Eds.) *Business process management workshops* (Vol. 132, pp. 256–267). Springer, Tallinn, Lecture Notes in Business Information Processing.
- Kwiatkowska, M., Milner, R., Sassone, V. (2004). Science for global ubiquitous computing. *Bulletin of the EATCS*, 82, 325–333.
- Latour, B. (1992). Where are the missing masses? The sociology of a few mundane artifacts. In Bijker, W.E., & Law, J. (Eds.) *Shaping technology/building society: studies in sociotechnical change* (pp. 151–180). Cambridge: MIT Press.
- Lindblom, J., & Ziemke, T. (2003). Social situatedness of natural and artificial intelligence: Vygotsky and beyond. *Adaptive Behavior*, 11(2), 79–96.
- Lupton, D. (2013). Understanding the human machine. *IEEE Technology and Society Magazine*, 32(4), 25–30.
- McBride, N. (2011). From social machine to social commodity: redefining the concept of social machine as a precursor to new Web development approaches. In *3rd international conference on web science*. Koblenz.
- Michelucci, P. (Ed.) (2013). *Handbook of human computation*. New York: Springer.
- Miller, M. (2015). *The Internet of Things: how smart TVs, smart cars, smart homes, and smart cities are changing the world*. Indianapolis: Que Publishing.
- Orlikowski, W.J., & Scott, S.V. (2008). Sociomateriality: challenging the separation of technology, work and organization. *The Academy of Management Annals*, 2(1), 433–474.
- Palermos, S.O. (2017). Social machines: a philosophical engineering. *Phenomenology and the Cognitive Sciences*, 16(5), 953–978.
- Pentland, A., & Liu, A. (1999). Modeling and prediction of human behavior. *Neural Computation*, 11(1), 229–242.
- Pezzulo, G., D'Ausilio, A., Gaggioli, A. (2016). Predictive technologies: can smart tools augment the brain's predictive abilities? *Frontiers in Neuroscience*, 10(186), 1–6.
- Piccinini, G. (2015a). Computation in physical systems. In Zalta, E.N. (Ed.) *The stanford encyclopedia of philosophy*, Summer 2015 Edn. Stanford: Stanford University.
- Piccinini, G. (2015b). *Physical computation: a mechanistic account*. Oxford: Oxford University Press.
- Ricci, A., Piunti, M., Tummolini, L., Castelfranchi, C. (2015). The mirror world: preparing for mixed-reality living. *IEEE Pervasive Computing*, 14(2), 60–63.
- Robertson, D., & Giunchiglia, F. (1987). Programming the social computer. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 371, 20120379.
- Rose, K., Eldridge, S., Chapin, L. (2015). *The internet of things: an overview*. Geneva: Tech. rep. Internet Society.
- Ruaro, M.E., Bonifazi, P., Torre, V. (2005). Toward the neurocomputer: image processing and pattern recognition with neuronal cultures. *IEEE Transactions on Biomedical Engineering*, 52(3), 371–383.
- Shadbolt, N., Van Kleek, M., Binns, R. (2016). The rise of social machines: the development of a human/digital ecosystem. *IEEE Consumer Electronics Magazine*, 5(2), 106–111.
- Shirky, C. (2010). *Cognitive surplus: creativity and generosity in a connected age*. London: Penguin Books Ltd.
- Simon, H.A. (1996). *The sciences of the artificial*, 3rd Edn. Cambridge: MIT Press.

- Smart, P.R. (2017a). Machine intelligence and the Social Web: how to get a cognitive upgrade. In Gripon, V., Chernavskaya, O., Smart, P.R., Primo, T.T. (Eds.) *9th international conference on advanced cognitive technologies and applications (COGNITIVE17)* (pp. 96–103). Athens: IARIA.
- Smart, P.R. (2017b). Situating machine intelligence within the cognitive ecology of the Internet. *Minds and Machines*, 27(2), 357–380.
- Smart, P.R. (2018). Human-extended machine cognition. *Cognitive Systems Research*, 49, 9–23.
- Smart, P.R., & Madaan, A. (2017). The social scaffolding of machine intelligence. *International Journal On Advances in Intelligent Systems*, 10(3&4), 261–279.
- Smart, P.R., & Shadbolt, N.R. (2014). Social machines. In Khosrow-Pour, M. (Ed.) *Encyclopedia of information science and technology*. Hershey: IGI Global.
- Sterelny, K. (2007). Social intelligence, human intelligence and niche construction. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1480), 719–730.
- Strohmaier, M., & Wagner, C. (2014). Computational social science for the World Wide Web. *IEEE Intelligent Systems*, 29(5), 84–88.
- Swan, M. (2013). The quantified self: fundamental disruption in big data science and biological discovery. *Big Data*, 1(2), 85–99.
- Taivalsaari, A., & Mikkonen, T. (2017). A roadmap to the programmable world: software challenges in the IoT era. *IEEE Software*, 34(1), 72–80.
- Theiner, G. (2014). Varieties of group cognition. In Shapiro, L.A. (Ed.) *The routledge handbook of embodied cognition* (pp. 347–357). New York: Routledge.
- Tromp, N., Hekkert, P., Verbeek, P.P. (2011). Design for socially responsible behavior: a classification of influence based on intended user experience. *Design Issues*, 27(3), 3–19.
- Turing, A.M. (1950). Computing machinery and intelligence. *Mind*, 59(236), 433–460.
- Walport, M. (2014). *The internet of things: making the most of the second digital revolution*. London: Tech. rep., UK Government Office for Science.
- Wilson, R.A., & Craver, C.F. (2007). Realization: metaphysical and scientific perspectives. In Thagard, P. (Ed.) *Philosophy of psychology and cognitive science* (pp. 81–104). Oxford: North-Holland.
- Ylikoski, P. (2018). Social mechanisms. In Glennan, S., & Illari, P.M. (Eds.) *The routledge handbook of mechanisms and mechanical philosophy* (pp. 401–412). New York: Routledge.
- Zednik, C. (2018). Mechanisms in cognitive science. In Glennan, S., & Illari, P.M. (Eds.) *The routledge handbook of mechanisms and mechanical philosophy* (pp. 389–400). New York: Routledge.