

Cyber-Physical Systems and Social Machines

Paul Smart

University of Southampton
Southampton, United Kingdom
ps02v@ecs.soton.ac.uk

Adrian Cox

University of Southampton
Southampton, United Kingdom
A.J.Cox@soton.ac.uk

Kieron O'Hara

University of Southampton
Southampton, United Kingdom
kmo@ecs.soton.ac.uk

Wendy Hall

University of Southampton
Southampton, United Kingdom
wh@ecs.soton.ac.uk

ABSTRACT

Social machines have long been a topic of theoretical and empirical interest for the Web and Internet science community. Although there is no widespread agreement as to the precise meaning of the term “social machine,” the term has served as something of a conceptual anchor for research efforts into a number of online systems, such as Facebook, Wikipedia, and Galaxy Zoo. Recently, the term “social machine” has surfaced in a somewhat different disciplinary context. In particular, it has been applied to systems that are commonly referred to as cyber-physical systems. Here, the residual ambiguity surrounding the meaning of the term “social machine” has stymied efforts to understand the nature of the relationship between social machines and cyber-physical systems. This situation is further complicated by the absence of a precise conceptual understanding of cyber-physical systems. The present paper outlines an approach to understanding both social machines and cyber-physical systems that is intended to resolve this impasse. Inspired by recent work in the philosophy of science, we present a mechanistic (or neo-mechanical) account of social machines that draws attention to the features of socio-technical mechanisms. Such an account, we suggest, is able to accommodate the intuitions of Web scientists regarding the nature of social machines, while simultaneously illuminating the nature of the relationship between social machines and cyber-physical systems. As an added bonus, the account speaks to the interdisciplinary concerns of the Web science community by highlighting a link between the science of social machines and mechanism-based approaches in the social sciences.

CCS CONCEPTS

• **Information systems** → **World Wide Web**; • **Human-centered computing** → **Collaborative and social computing theory, concepts and paradigms**; *Ubiquitous and mobile computing*; • **Networks** → *Cyber-physical networks*.

KEYWORDS

Social Machines, Internet of Things, Cyber-Physical Systems, Mechanisms, Socio-Technical Systems, Analytical Sociology

1 INTRODUCTION

Social machines are a prominent focus of research attention for the sciences of the Web and Internet [27, 28, 42]. For the most part, the term “social machine” has been applied to a familiar array

of Web-based systems. These include systems such as Wikipedia [22, 27, 37, 41], Facebook [27], Twitter [22], and Galaxy Zoo [22, 41]. Recently, however, some have called for an expansion in the scope of the social machine research effort. De Roure et al. [14], for example, suggest that the “concept of Social Machines has become an established lens to describe the sociotechnical systems of Web Science” and that such a lens can be applied to the study of Cyber-Physical Systems (CPSs) [14, p. 65].

The primary aim of the present paper is to explicate the nature of the relationship between social machines and CPSs. Our understanding of this relationship is complicated by the absence of precise definitions for both the social machine and CPS concepts. From an analytical perspective, then, we confront two challenges. The first is a normative conceptual challenge: How ought we to conceptualize social machines and CPSs? The second challenge is one of understanding the nature of the relationship between social machines and CPSs. Is it the case, for example, that some CPSs can be regarded as social machines? Or is the relationship between social machines and CPSs somewhat more methodological in nature (e.g., CPSs ought not to be regarded as social machines, but some of the methods developed to study social machines are applicable to CPSs)?

In the present paper, we outline a theoretical account of social machines that is intended to serve as the conceptual foundation for future theoretical and empirical work. This account is what we call the *mechanistic view* of social machines. We show how the mechanistic view is able to accommodate many of the intuitions voiced by social machine researchers about the nature of social machines. We also map out the nature of the relationship between social machines and CPSs with a view to facilitating interdisciplinary collaboration.

The structure of the paper is as follows: Section 2 provides an overview of the mechanistic view of social machines. Our aim in this section is to provide a summary of the main concepts associated with the mechanistic view; in particular, we focus our attention on those concepts that are relevant to the remainder of the paper. Section 3 provides an overview of CPSs. Here, we review some of the definitions that have been offered in respect of CPSs. We also describe some of the ways of thinking about CPSs. Section 4 extends the discussion of CPSs to include what are called Cyber-Physical-Social Systems (CPSSs). These systems, we suggest, provide a point of contact with the notion of a social machine, which is established courtesy of the mechanistic view. In Section 5, we highlight how a mechanistic approach to social machines provides the conceptual foundation for interdisciplinary work that spans the social,

computational, and engineering sciences. Section 6 concludes the paper.

2 SOCIAL MACHINES

Social machines were first introduced into Web science by Berners-Lee and Fischetti [4]. According to their characterization:

Real life is and must be full of all kinds of social constraint—the very processes from which society arises. Computers can help if we use them to create abstract social machines on the Web: processes in which the people do the creative work and the machine does the administration. [4, p. 172]

This is what we will call the *content creation view* of social machines. According to this view, we confront a social machine whenever we encounter a process in which there is a discernible division of labor between the human and technological elements of a Web-based system. In particular, the contributions of the human participants should correspond to a form of creative activity, whereas the contributions of the technological elements should correspond to a form of administrative activity.

In recent years, there have been a number of further attempts to characterize social machines. Hooper et al. [29], for example, conceive of social machines as problem-solving systems, i.e., socio-technical systems that aim to resolve problems in a socially distributed fashion. A not altogether incompatible view is proposed by Palermos [37]. Palermos suggests that we should regard social machines as distributed cognitive systems, i.e., as systems in which cognitive processing routines (e.g., problem-solving processes) are distributed across human agents and technological artifacts. Social machines have also been approached from a socio-computational perspective. Horsman [30], for example, characterizes social machines as “computational ecosystems comprising both digital computers and multiple human users, all acting towards a computational goal” [30, p. 2].

The view of social machines to be outlined here is what is known as the *mechanistic view* [see 44]. From the standpoint of the mechanistic view, social machines are online systems whose events, states, and processes are realized by a socio-technical mechanism—that is, a mechanism consisting of both social components (e.g., human individuals) and technological components (e.g., conventional digital computers). The idea, in essence, is that we discern a social machine whenever our attempts to explain some form of online phenomenon lead us to the conclusion that a socio-technical mechanism is responsible for the phenomenon in question.

The mechanistic view of social machines is rooted in a body of philosophical work that has come to be known as *mechanical philosophy* [19, 20]. This is a specialist area of the philosophy of science that concerns itself with the nature of mechanisms and the role they play in our scientific understanding of the world [10]. It should come as no surprise, therefore, to learn that mechanisms lie at the heart of the mechanistic view of social machines. While there are disagreements as to the precise meaning of the term “mechanism,” the central features of the mechanism concept are captured in the following generic definition by Illari and Williamson [32]:

A mechanism for a phenomenon consists of entities and activities organized in such a way that they are responsible for the phenomenon. [32, p. 120]

As is clear from this definition, the building blocks of mechanisms are what are called “entities” and “activities.” Entities are typically conceived as material objects, along with their associated properties. They are the physical parts of mechanisms—the things that make up the mechanism. Activities, by contrast, are typically conceptualized as the “producers of change” [35, p. 3] and as the “causal components of mechanisms” [9, p. 6]. They “are the things that the entities do” [12, p. 16]. In many cases, the activities describe the nature of the interactions between the entities that make up a mechanism, as when we say that an enzyme (entity) phosphorylates (activity) a protein (entity), a neuron (entity) releases (activity) a neurotransmitter (entity), a human agent (entity) edits (activity) a Wikipedia entry (entity), and a human agent (entity) tags (activity) an online image (entity).

Illari and Williamson’s [32] definition is the one we will adopt in the present paper. Accordingly, we will define a socio-technical mechanism as a mechanism that consists of entities of a particular type. In particular, we suggest that a socio-technical mechanism is a mechanism that consists of entities of two types. These are what we will call *social components* and *technological components*. For the purposes of this paper, we will assume that the social components are human individuals and the technological components are digital computational systems or devices. A socio-technical mechanism is thus a mechanism in which human individuals (*qua* social components) and digital computers (*qua* technological components) interact in such a way as to be responsible for some kind of phenomenon.

Figure 1a shows a common depiction of a mechanism, which is attributable to Craver [9]. Craver refers to the mechanism as S and the entities of the mechanism as X_1, X_2, \dots, X_n . S ’s activity is denoted by the symbol Ψ . S ’s Ψ -ing is deemed to represent the behavior of the mechanism as a whole. The activities of the entities that comprise the mechanism are denoted by $\Phi_1, \Phi_2, \dots, \Phi_n$. Together, the entities and activities are what are called the *components* of the mechanism. The directed arrows represent the causal interactions between the components. It is these causal interactions that are deemed to be responsible for S ’s Ψ -ing.

Figure 1b depicts a socio-technical mechanism using Craver’s diagrammatic scheme. This is intended to represent the sort of mechanism that is targeted by the mechanistic view of social machines. The first thing to note about this mechanism is that it consists of multiple human individuals. We assume that in order for a mechanism to count as a socio-technical mechanism, it must consist of *multiple* human individuals. This is what we will refer to as the *sociality criterion*. The sociality criterion is important, because it is at best unclear that a mechanism consisting of a single human individual ought to be counted as a *socio-technical* mechanism. Socio-technical mechanisms, we suggest, are best seen as a particular kind of *social mechanism*, which are the sort of mechanisms that are the target of mechanism-oriented approaches in the social sciences [e.g., 26, 52]. In this sense, the mechanistic view provides a conceptual foundation for interdisciplinary work with the social sciences, especially with approaches that seek to explain social phenomena via an appeal to social mechanisms [e.g., 24] (see Section 5, for more on this).

In addition to multiple human individuals, a socio-technical mechanism consists of one or more technological components.




Figure 1: A graphical depiction of a phenomenon (top) and its mechanism (bottom) [adapted from 9]. (a) In a general sense, a mechanism consists of components (X_1, X_2, \dots, X_n) that interact (arrows) in such a way as to constitute or realize the phenomenon of interest. (b) A socio-technical mechanism of the sort targeted by the mechanistic view.

For present purposes, we can think of technological components as conventional digital computers, although nothing about the mechanistic view precludes the idea that other kinds of technology could form part of a socio-technical mechanism.¹

In order for a mechanism to qualify as a socio-technical mechanism it must, therefore, consist of multiple social components and one or more technological components. What makes these components part of a common mechanism is the fact that they are deemed constitutively relevant to an explanandum phenomenon (i.e., the phenomenon that is the target of a mechanistic explanation). There is an important distinction to be made here between the notion of *causal relevance* and the notion of *constitutive relevance* [see 8]. Causal relevance identifies the material elements that are causally relevant to a phenomenon, in the sense of helping to produce, generate, or bring about the phenomenon. Constitutive relevance, by contrast, identifies the components that realize or constitute a phenomenon.² From the standpoint of the mechanistic view, the kind of relevance we are looking for is constitutive relevance. It is thus not enough to say that a particular phenomenon was caused by some set of social and technological elements. Instead, what we are looking for is a state-of-affairs in which the social and technological elements form part of the physical fabric that realizes or constitutes a particular phenomenon. In most cases, this means that the social and technological elements must work together as an integrated whole, with individual components interacting with other components, pretty much as the components of an internal combustion engine work together to realize the locomotory phenomena of a conventional automobile. In all likelihood, it is this requirement for *socio-technical integration* that underlies

a number of recent attempts to define social machines. Within the social machine literature, we thus encounter the idea of social machines as “purposefully designed sociotechnical system[s] comprising machines and people” [15], as systems in which “the human and digital parts... [form] a machine in which the two aspects are seamlessly interwoven” [43], as “large-scale sociotechnical systems that combine human and machine computations and interactions into a new system with emergent characteristics” [1], as systems in which the “components... are people and technologies” [36], as systems that involve “the co-constitutional involvement of humans and technologies” [46], as systems whose “behaviour [is] co-created by human participants and technological components” [14], and as systems that are constituted by “humans, machines, and their interactions” [45].

The arrows in Figure 1b represent the causal links between the components. These arrows indicate the diversity of interactions that are permissible under the mechanistic view. Note, for example, that component ‘interactions’ can be either unidirectional or bidirectional. Similarly, the interactions can be of various kinds. We thus have examples of human–human, machine–machine, and human–machine interactions. All of these interactions are permissible under the mechanistic view. The only constraint imposed on the causal exchanges is that *some* of these exchanges must involve the use of Internet protocols. This is what we will call the *Internet Protocol (IP) criterion*. The IP criterion is intended to restrict the notion of a social machine to the Web or Internet. In short, the IP criterion is intended to focus our attention on a specific class of socio-technical systems, namely, those that qualify as Web- or Internet-based systems (or, more informally, online systems). This is consistent with the scope of the social machine research effort. For the most part, research into social machines focuses on systems that rely (at least in part) on the use of Web-specific application protocols, such as HyperText Transfer Protocol (HTTP). Examples of such systems include Facebook, Wikipedia, Twitter, Galaxy Zoo, and so on. In

¹It should also be noted that social and technological components need not exhaust the constituents of a socio-technical mechanism. That is to say, a socio-technical mechanism may include components that do not qualify as social or technological.

²Note that in the present paper we treat the notions of *mechanistic constitution* [e.g., 3] and *mechanistic realization* [e.g., 50] as synonymous.

addition to these Web-based systems, it has been suggested that the social machine concept is also applicable to Internet-based systems [1, 15, 29, 44]. Such claims are typically made in relation to the communication protocols used by mobile apps and games, as well as the protocols employed to support the Internet of Things (IoT).

To summarize the mechanistic view: Social machines are systems that are associated with one or more socio-technical mechanisms, where the notion of a socio-technical mechanism is to be understood as a particular kind of (hybrid) mechanism that is composed of elements drawn from both the social (e.g., human individuals) and technological (e.g., conventional digital computers) realms. What makes something a social machine is the fact that a socio-technical mechanism realizes the phenomena exhibited by the larger system (i.e., the social machine). From the standpoint of the mechanistic view, then, the presence of a socio-technical mechanism is used to determine whether some online system ought to be counted as a *bona fide* social machine. The core idea is that when we encounter some system on the Web—Wikipedia, let's say—we should judge whether this system is a social machine based on the presence of a socio-technical mechanism that constitutes or realizes one of the phenomena associated with the relevant system.

3 CYBER-PHYSICAL SYSTEMS

The traditional focus of the social machine research effort includes systems such as Facebook, Wikipedia, Galaxy Zoo, and so on. These are what we might call traditional or classic examples of social machines. In recent years, however, the term has been applied to systems of a seemingly different nature. One example of this stems from work by De Roure et al. [14]. In their paper, “Towards a Cyberphysical Web Science,” De Roure et al. discuss the application of the term “social machine” to the augmented reality game *Pokémon Go!* The details of this work need not concern us here. What matters, for present purposes, is simply the idea that there exists some sort of relationship between the notion of a social machine and a CPS. This relationship is deemed to be of interest to Web scientists, De Roure et al. suggest, on the grounds that “Web Science is becoming ‘cyberphysical’, with a deepening intertwining of the digital and physical worlds and the algorithms that connect them” [14, p. 65].

In order to understand the nature of the relationship between social machines and CPSs we need to understand what is meant by the terms “social machine” and “cyber-physical system.” The mechanistic view provides us with a theoretical account of what it means for something to count as a social machine, but it clearly tells us nothing about what it means for something to count as a CPS. This is problematic because there is no widely-accepted definition of the term “cyber-physical system.” Some of the definitions of CPSs on offer are the following:

The concept of a cyber-physical system is a generalization of embedded systems.³ A cyber-physical system consists of a collection of computing devices communicating with one another and interacting with the physical world via sensors and actuators in a feedback loop. [2, p. 1]

Cyber-Physical Systems (CPS) are integrations of computation and physical processes. Embedded computers and networks monitor and control the physical processes, usually with feedback loops where physical processes affect computations and vice versa. [34, p. 363]

Cyber-physical systems (CPS) are physical and engineered systems whose operations are monitored, coordinated, controlled and integrated by a computing and communication core. [40, p. 731]

A cyber-physical system (CPS) is composed of a physical system and its corresponding cyber systems that are tightly fused at all scales and levels. [51, p. 320]

Cyber-physical systems (CPSs) enable the physical world to merge with the virtual world by integrating computation and physical processes. A CPS facilitates tight integration between computation, communication, and control in its operation and interactions with the environment in which it's deployed. [21, p. 4]

The problem with these definitions is that they tend to raise more questions than they answer. What does it mean, for example, for computational processes to be integrated with physical processes? What exactly is a cyber system? And how is a cyber system distinct from a physical system, given that a computational system (e.g., a conventional computer) would seem to qualify as a *bona fide* physical system [see 38]?

Clearly, we will not have the space to address all these issues in the present paper. With this in mind, we will assume that the term “cyber system” is a reference to some form of digital computational device (e.g., a conventional digital computer or a microcontroller). We will then assume that the term “physical system” is a reference to any system that does not qualify as a cyber system. Based on these assumptions, the term “cyber-physical system” is deemed to refer to any system that involves some form of interaction between a cyber system and a physical system. In essence, a CPS is a hybrid system that features the exchange of information, energy, and/or materials between a (digital) computational system and a system that is not typically thought of in computational terms. This does not mean that the larger systemic organization (i.e., the hybrid system) should, itself, be thought of as a computational system (although nothing discounts that possibility). A washing machine, for example, is not, by itself, typically thought of as a computational system; nevertheless, its behavior is, at least in part, the product of computational processes.

There are, in fact, a number of ways in which this appeal to hybridity can be understood. These include the following:

- **Compositional View:** A CPS is a system that is composed of both a cyber system and a physical system.
- **Controller View:** A CPS is a system in which a cyber system controls a physical system.
- **Constitutional View:** A CPS is a system whose phenomena (i.e., events, states, and processes) are realized by a hybrid mechanism, which itself consists of a cyber system and a physical system (i.e., the cyber and physical [sub]systems are constitutively relevant to the phenomena exhibited by the larger [hybrid] system).

³This is the view of embedded systems that we will adopt in the present paper. We thus suggest that the class of embedded systems ought to be seen as a proper subset of the class of CPSs. It follows, therefore, that every embedded system will qualify as a CPS.

Of the three views on offer, it should be clear that the constitutional view is the one that is most closely aligned with the mechanistic view of social machines described in Section 2. That being said, it should also be clear that the sort of mechanisms targeted by the constitutional view are not the same as those associated with the mechanistic view. The mechanistic view, recall, requires us to identify mechanisms that are constituted (*inter alia*) by social and technological components. It also mandates that some of the operations of these mechanisms will be mediated via the use of Internet protocols. Neither of these features are explicitly required by the constitutional view, although it is also true (and this is the important bit) that there is nothing about the constitutional view that would rule out the possibility of some CPSs qualifying as social machines. For a CPS to qualify as a social machine under the constitutional view, the CPS will need to feature mechanisms that include (as components) multiple human individuals (the sociality criterion) and digital computational systems/devices. It will also need to feature mechanisms that rely on the use of Internet protocols (the IP criterion). As we will see, neither of these constraints pose much of a problem for the constitutional view of CPSs. In respect of the IP criterion, for example, it is perfectly permissible for CPSs to exploit IoT or Web of Things (WoT) devices as part of their functional operation, and we have already seen that there is nothing about the mechanistic view that would preclude the use of such devices in a social machine context. There is, of course, nothing about the constitutional view of CPSs that would *mandate* the use of IoT devices as part of a socio-technical mechanism; nevertheless, much of the research into CPSs is centered around the use of such devices [6, 13]. In particular, IoT devices are often used as part of the sensor and actuator suite of contemporary CPSs. There is also considerable interest in understanding how the Internet and Web can be exploited to support the *ad hoc* assembly of different CPSs based on a common commitment to Internet- and/or Web-related standards. Such standards include those governing network-mediated communication (e.g., IP-compatible networks), service provision (e.g., Web services), and data representation (e.g., the use of Semantic Web formalisms).

Relative to the claim that some CPSs may qualify as social machines, then, the IP criterion is not particularly problematic: there appears little reason to doubt that (at least some) CPSs will exploit the communicative and representational resources of the Internet and Web as part of their functional operation. Much the same can be said about the sociality criterion, although we reserve a discussion of this issue for Section 4. For the time being, it is worth noting that the constitutional view should not be viewed as an *ad hoc* characterization of CPSs—one that is intended to service the interests of the mechanistic view. That is to say, we should not see the constitutional view as a form of definitional contrivance. Instead, the constitutional view is intended to reflect the definitional efforts of those in the CPS community. Consider, for example, the way in which Rajkumar et al. [39] characterize CPSs:

The National Science Foundation defines cyber-physical systems (CPS) as “engineered systems that are built from, and depend upon, the seamless integration of computational algorithms and physical components”—that is, cyber

and physical components. In practical terms, this integration means that, to understand CPS behavior, we cannot focus only on the cyber part or only on the physical part. Instead, we need to consider both parts working together. [39, p. xiii]

What is interesting about this characterization is the emphasis placed on integration and understanding. This echoes one of the central themes of the mechanistic view—the idea that in order to understand or explain some system-level phenomenon we are required to advert to an explanatory account that identifies and describes the mechanism that is responsible for the explanandum phenomenon. Obviously, Rajkumar et al. [39] are not relying on the phraseology adopted by those who analyze mechanism-related concepts. Nevertheless, their definition does appeal to ideas that are prevalent in the relevant philosophical literature. This includes the idea that explanatory efforts are guided by a consideration of “parts” that are, in some sense, “working together.” Such characterizations echo the sort of claims made by the proponents of mechanical philosophy. Craver [11], for example, refers to mechanism components as “working parts.” Components, he suggests, “are not mere spatial or temporal parts but working parts—parts that are involved in, contribute to, or make a difference to the property or activity of the whole” [11, p. 14].

4 INCORPORATING THE SOCIAL

In Section 2, we saw that in order to confirm the status of a system as a social machine we need to isolate a mechanism that satisfies the sociality criterion. What this means is that a mechanism should consist of two or more human individuals. As things stand, there is nothing about the notion of a CPS that would suggest that some CPSs ought to be regarded as social machines. This is because there is no reason to think that any CPSs are required to satisfy the sociality criterion. In recent years, however, there has been a growing interest in a specific class of CPSs that does make explicit reference to the social realm. These systems are what are known as Cyber-Physical-Social Systems (CPSSs). Examples of such systems are many and varied. They include the likes of smart cities [5, 7, 13], smart enterprises [49], smart transportation systems [51], environment monitoring systems [13, 17], vehicle platooning systems [16], and smart senior caring systems [31].

As with CPSs, there is no widely accepted definition of the term “cyber-physical-social system.” As a specific kind of CPS, we can assume that CPSSs inherit the features of CPSs. Similarly, it is at least plausible that all of the views discussed in Section 3 (i.e., the compositional, controller, and constitutional views) are applicable to CPSSs. As with the more generic notion of a CPS, it is the constitutional view that is the most closely aligned with the mechanistic view of social machines. According to a constitutional view of CPSSs, a CPSS is a system whose phenomena (i.e., events, states, and processes) are realized by a mechanism that consists of components that are drawn from the physical, social, and cyber realms.

At this point, we are in a position to explicate the relationship between social machines and CPSSs. From the standpoint of a constitutional perspective, CPSSs will satisfy the sociality criterion and they are thus a suitable candidate for inclusion in the class




Figure 2: The class of social machines occupies a region that lies at the intersection of cyber systems and social systems (shaded ellipse). [CPS: Cyber-Physical System; PSS: Physical-Social System; CSS: Cyber-Social System; CPSS: Cyber-Physical-Social System.]

of social machines. Providing such systems also meet the IP criterion (e.g., they involve the use of IoT devices), then such systems will, we suggest, qualify as social machines. The nature of this relationship is depicted in Figure 2. In this figure, the class of social machines is represented by the shaded ellipse. Social machines are thus seen to occupy a region that lies at the intersection of cyber systems (i.e., computational systems) and social systems. In effect, social machines are conceptualized as hybrid systems that straddle the computational and social realms. A particular subset of these systems will also qualify as CPSSs (and thus CPSs). Such systems lie at the intersection of physical systems, cyber systems, and social systems. For the sake of convenience, let us refer to this particular class of social machines as Cyber-Physical Social Machines (CPSMs). Such systems are importantly distinct from social machines of the more traditional variety, such as Wikipedia or Galaxy Zoo. They are distinct on account of the fact that CPSMs qualify as *bona fide* members of the class of CPSs. This is not the case for systems that have been the typical focus of attention for the social machine community. A system such as Galaxy Zoo, for example, will qualify as a Cyber-Social System (CSS), but there is no reason to think that it ought to be recognized as a CPSS. According to the conceptual approach on offer, a social machine will only qualify as a CPSM if it also qualifies as a CPS. When this is the case, it will also qualify, we suggest, as a CPSS. In this sense, all the examples of CPSSs mentioned in the scientific and engineering literature (e.g., smart cities, smart transportation systems, and so on) are at least candidate members of the class of social machines (specifically, the class of CPSMs).

The mechanistic view of social machines thus provides the basis for a particular way of thinking about the relationship between

social machines and CPSs. This is no doubt useful when it comes to our conceptual understanding of CPSs and social machines, but the ‘market value’ of the mechanistic view is likely to hinge on more than just an improvement in how we think about social machines. In all likelihood, there are additional issues to be addressed before the conceptual approach on offer ought to be regarded as of any practical significance or relevance. For a start, it is not particularly clear that a commitment to the constitutional view really amounts to much. What difference does it really make if we view CPSSs from the standpoint of the controller view as opposed to the constitutional view? Does this kind of conceptual distinction really mark a significant shift in the practical effort to study, understand, evaluate, and (importantly) design CPSSs? Similarly, it is not particularly clear why we should care about the distinction between CPSMs and social machines of the more traditional variety. Does this conceptual distinction really amount to anything of practical interest or relevance for the Web and Internet science community? Does it, for example, alter the way we study, model, and build social machines?

These are important issues, and they ought to be given a far more detailed airing than we can offer here. For present purposes, however, will attempt to provide something in the way of an initial response to these issues.

First, let us consider the relative merits of the constitutional view, as opposed to the controller view. The main advantage of the constitutional view is that it encourages us to consider the role of the human social environment in securing certain system-level properties. From the standpoint of the constitutional view, we confront an integrated socio-technical system, and the properties of this system are related to the functional inter-operation between the social and technological constituents of the system. This means that properties like resilience, reliability, security, energy efficiency, trustworthiness, and so on, might supervene on mechanisms that factor in the contributions of human participants.

As a means of helping us get a better grip on this idea, consider the way in which we might approach the phenomenon of human hand movements from the standpoint of the controller view. The human hand is evidently distinct from the biological brain, at least in an anatomical sense, and it seems perfectly legitimate to regard the brain as a source of control signals for specific hand movement trajectories. From the standpoint of the controller view, then, we will be inclined to view the biological brain as a form of control system. The hand, in this case, is not part of the material mix that yields a solution to the relevant control problem; it is instead, part of the problem that must be solved. Presumably, then, the mechanisms that yield a solution to this problem will not lie within the hand itself. They must lie elsewhere; in all likelihood, they are neural mechanisms located within the biological brain.

Now contrast this with the perspective afforded by the constitutional view. From the standpoint of the constitutional view, we will see the hand not so much as something that lies external to the control mechanism but as an intrinsic part of the mechanism that solves the control problem.

Interestingly, of these two perspectives, it is the constitutional perspective that appears to be vindicated by empirical research. Using a combination of real-world cadaveric experiments and computer simulations, Valero-Cuevas et al. [48] show how the tendon

network of the human hand forms part of an anatomically distributed information processing system for the control of finger movements. The mechanisms responsible for movement control are thus *not* contained solely within the biological brain; instead, they are distributed across the nervous system and tendon network, such that “part of the controller is embedded in the anatomy, contrary to current thinking that attributes the control of human anatomy exclusively to the nervous system” [48, p. 1165].

The shift in perspective (from the controller view to the constitutional view) is important, for it alters the way we think about the design of CPSMs. In particular, a constitutional approach encourages us to consider the various ways in which the human social environment may play a productive role in solving particular problems, or at least influencing the kinds of things we regard as relevant to the problem-solving effort. Such a perspective is already arguably well-represented by existing work in Web science, especially that which focuses attention on the way in which human individuals (or collections thereof) are *incorporated* into computational routines [e.g., 33]; it is, however, a perspective that remains to be fully explored in a cyber-physical setting.

Let us now direct our attention to the second issue raised above: the issue about the merits of distinguishing between traditional social machines and CPSMs. In respect of this issue, it is important to consider some of the features that make CPSs worthy of specialist scientific and engineering attention. All CPSMs, recall, are CPSSs (and thus CPSs), so the features that make CPSSs deserving of specialist scientific/engineering attention will also be those that make CPSMs worthy of specialist scientific/engineering attention. One such feature concerns the particular technical challenges associated with the implementation of CPSs. Such challenges are well-recognized. They include issues of concurrency, reactive computation, real-time computation, and the challenges of integrating discrete-time and continuous-time systems [see 2, pp. 2–5]. Other challenges relate to issues of (big) data processing, mobility, and the need for decentralized control schemes. Such challenges are typically seen to require a shift in the methods and practices of computer science.

In addition to the technical challenges, there are reasons to think that CPSMs ought to be distinguished from traditional social machines on account of their socially-relevant properties. Consider, for example, that many traditional social machines are founded on voluntary forms of social participation. There is, for example, no reason why individual human agents are obliged to participate in a citizen science system, such as Galaxy Zoo. This means that the designers of traditional social machines need to give due care and attention to the forces and factors that motivate human participation in the system [e.g., 47].

For the most part, then, traditional social machines (e.g., Facebook, Twitter, Wikipedia, Galaxy Zoo, and so on) are what we might call *facultative social machines*. They are social machines that feature voluntary forms of social participation. In contrast to facultative social machines, *obligate social machines* are social machines characterized by more obligatory forms of social participation. An example of such a system might be a smart transportation system, where one’s participation in the system is not so much a matter of choice, as it is a *de facto* feature of social life: if an individual wants to travel from A to B, they might have little choice but to participate

in the system. One could, of course, opt to refrain from travelling, but this might require a more general retreat from society. And with the transition to a smart society (and, perhaps ultimately, a smart planet), there are likely to be progressively fewer and fewer places where one can retreat to.

This, then, is one of the reasons why the distinction between traditional social machines and CPSMs matters. In particular, we suggest that traditional social machines are more likely to be of the facultative rather than the obligate variety, while the reverse is likely to be true for CPSMs. This distinction is not clear-cut, of course, since some CPSMs may be more facultative than others. Nevertheless, the distinction is important when it comes to understanding some of the features of CPSMs. Compared to traditional social machines, CPSMs are more likely to be of the obligate variety and this necessitates a careful consideration of their social impact (as well as, perhaps, the ethics of their deployment). For traditional social machines issues of incentivization and motivation are all-important; for CPSMs, the focus is more likely to be on issues of social acceptance and inclusivity. Traditional social machines are like places we visit, but CPSMs may ultimately become the place we call our social home.

Additional reasons to merit a distinction between traditional social machines and CPSMs relate to issues of privacy, safety, security, and economic impact. A temporary disruption in Facebook’s service offering is no doubt annoying, but a similar disruption to a smart city application or a smart transportation system is apt to be much more profound. If nothing else, the financial costs associated with these disruptions are not the same: the costs of disrupting Facebook are mostly felt by a corporate entity; the costs associated with the disruption of a smart transportation system are likely to be felt across multiple sectors of society.

This is not to say that disruptions to traditional social machines are inconsequential or unimportant, but it is surely the case that the shift to CPSMs does entail an accentuation of the sort of concerns (e.g., privacy, security, etc.) that accompany their traditional counterparts. A smart city application or smart transportation system is not just an application that is tacked onto the social life of the city, it is (especially according to the constitutional view) an integral part of what makes a (smart) city the thing it is. The claim, then, is that CPSMs are apt to feature an accentuation of the sort of concerns that are already associated with traditional social machines. This has implications for how CPSMs are designed, modelled, and developed. In particular, it is likely that CPSMs will need to comply with the sort of constraints that are typically applied to CPSs, including, for example, a susceptibility to formal verification methods [see 39, chap. 6].

While there are reasons to see CPSMs as distinct from traditional social machines, it is also important that we do not overplay such differences. In particular, it is probably a mistake to see CPSMs as wholly distinct from traditional social machines. This is one of the *disadvantages* of the conceptual scheme depicted in Figure 2, since it encourages us to think of CPSMs as entirely distinct from traditional social machines. In reality the picture is much more complicated, since there is no reason why a traditional social machine should not be incorporated into a larger system that, itself, qualifies as a CPSM. As noted by De et al. [13], there are a number of ways in which social media systems might be exploited in the context of CPSSs. A

specific example of this comes from a study by Du et al. [17]. They describe how the analysis of Twitter messages might be used in the context of an air quality monitoring system. In particular, they suggest that social media data can be used in conjunction with air quality sensors to support air quality assessments. The advantage of this approach is that it affords an opportunity to gather information about the subjective experiences and reactions of urban residents in a manner that is difficult (if not impossible) to acquire from conventional (technological) sensors.

It is also worth noting that nothing about the concept of a CPSM, as it is presented here, precludes the use of technologies that are typically encountered in the case of traditional social machines. One example of how a familiar array of Web technologies might be incorporated into a CPSM stems from work by Costanzo et al. [7]. They describe a smart city application, called Wi-City-Plus, which combines the use of social data with Semantic Web technology to provide a range of information services to urban residents. As part of their approach, Costanzo et al. describe how various bodies of data are mapped to a common ontology and then processed using rules expressed in a query language developed specifically for the Semantic Web, namely, SPARQL.

5 SOCIAL MECHANISMS AND SYNTHETIC SOCIOLOGY

Web science has long been recognized as a research effort that invites (and perhaps requires) productive forms of engagement and collaboration with a rich array of academic disciplines. Given that the Web is typically conceived as a socio-technical system, it is relatively easy to identify some of the disciplines that might be relevant to Web science (computer science and social science being two of the more obvious candidates). It is, however, much harder to specify *why* such interdisciplinary efforts are required, or what the common conceptual ground for such interdisciplinary efforts might be. Courtesy of the emphasis placed on socio-technical mechanisms, the mechanistic view of social machines provides us with a potential answer to such questions. In particular, the mechanistic view establishes contact with mechanism-based approaches in the social sciences, i.e., approaches that seek to explain social phenomena by detailing the mechanisms responsible for such phenomena.

Within contemporary sociology, the mechanistic approach finds its clearest and most explicit expression in what has come to be known as *analytical sociology* [23, 24]. This is a sub-field of contemporary sociology that focuses its attention on the mechanisms that underlie social phenomena. Such mechanisms, it is claimed, consist of individuals and the relationships that exist between individuals:

Analytical sociology explains by detailing the mechanisms through which social facts are brought about, and these mechanisms invariably refer to individuals' actions and the relations that link actors to one another. [25, p. 4]

The emphasis that analytical sociology places on social mechanisms is perfectly consistent with the mechanistic view of social machines. The only substantial difference is that the mechanistic view encourages us to see the technological elements of the Web and Internet as constitutively relevant to the sort of phenomena that are the focus of explanatory efforts in the social sciences. In essence, the mechanistic view encourages us to see the Web and

Internet as part of the machinery of the human social world—part of the causally-active physical fabric that makes certain types of social phenomena materially possible.

The extent to which the proponents of analytical sociology might be willing to embrace this idea is unclear, since much of the work in analytical sociology expresses a commitment to individualist doctrines that emphasize the explanatory importance of human individuals, sometimes to the exclusion of other entities [see 26]. There is, of course, a sense in which analytical sociologists are right to emphasize the explanatory importance of human individuals when it comes to the explanation of social phenomena. As noted by Hédstrom [23]:

...if we were able to press a pause button that suddenly froze all individuals and prevented them from performing any further actions. All social processes would then come to an immediate halt. (p. 28)

The mistake, however, is to assume that just because human individuals are important (perhaps essential) components of social mechanisms that *all* social mechanisms must be constituted by human individuals to the exclusion of other entities. The danger with such an approach is that it risks blinding us to the possibility that certain kinds of social phenomena might be realized by mechanisms that consist of more than just human agents. This is not to say that social phenomena can be realized in the absence of (multiple) human individuals, but it is surely the case that the Internet and Web are at least candidates for inclusion into social mechanisms. After all, if we tweak the aforementioned counterfactual posed by Hédstrom and ask ourselves what would happen if we were to press a button that suddenly disabled all the computers on the Internet and the Web, then it is not entirely clear that our contemporary society could continue to operate in quite the same way as it does at the moment. No doubt *some* social processes would survive the technological cataclysm, but a great many processes would also stand to be changed, or perhaps suspended altogether.

There is, no doubt, much more to be said about the extent to which the technological elements of the Web and Internet ought to be included in social mechanisms (and thus referenced as part of mechanistically-oriented social explanations). For present purposes, however, it is at least plausible that certain kinds of social phenomena are realized by the sort of mechanisms that lie at the heart of the mechanistic view of social machines. In this sense, the mechanistic view provides the conceptual foundation for interdisciplinary work between the science of social machines and mechanism-based approaches in the social sciences. Such work centers on the need to discover and describe the mechanisms responsible for social phenomena. At present many of these phenomena are associated with systems that are situated in the 'online' (or virtual) world (e.g., Facebook). With the advent of CPSMs, however, we anticipate a shift in the focus of explanatory efforts towards systems that are more typically seen as being situated in the 'offline' (or real) world (e.g., smart cities). In both cases, however, it seems likely that our attempts to explain a rich array of social phenomena will need to allow for the possibility that such phenomena are realized by socio-technical mechanisms.

Explanatory concerns are, however, not the only basis for interdisciplinary collaboration. From the standpoint of the mechanistic

view, the Web and Internet are not just factors that shape the human social world, they are part of the physical machinery that makes the human social world the thing it is—socio-technical mechanisms are, if you like, the “nuts and bolts” of society [see 18], the material building blocks of our social reality. This is important, for it depicts the development of social machines as a form of social engineering—one that focuses attention on the development of mechanisms that realize social phenomena. This, we suggest, provides the basis for a synthetically-oriented shift in analytical sociology: from an analytic effort directed at mechanism discovery towards a synthetic effort directed at mechanism implementation. We thus envisage a synthetic counterpart to analytical sociology—let us call it *synthetic sociology*—that focuses on the way in which the development of social machines is apt to change society. By itself, of course, this proposal is not particularly novel, since no one is likely to dispute the idea that the Web and Internet are capable of effecting shifts in the social landscape. What is new about the present proposal is the idea that social machines serve as the realization base for social phenomena. By building social machines, we are not just *enabling* new social processes to emerge, nor are we simply engaged in an effort that *allows* certain forms of social change to occur. Instead the effort to build social machines is one that provides an opportunity to forge *new* kinds of social mechanisms. And by building such mechanisms, we are creating the realization base for potentially *new* kinds of social phenomena. In this sense, social machines are not just an enabler for social processes, for in building a social machine we are creating the physical substrate that makes certain kinds of social phenomena materially possible. Similarly, the implementation of new social machines should not be seen solely as a means of enabling social change, for in implementing social machines we are (quite literally) *reconstituting* and thus reconstructing the social world, i.e., creating the physical mechanisms that are the material underpinnings of our social reality.

A commitment to synthetic sociology thus entails a subtle shift in the ‘direction’ of analytical sociology. For the most part, analytical sociology is concerned with the attempt to understand extant social phenomena, and such understanding inheres in the identification and description of mechanisms that are deemed to be responsible for these phenomena. Synthetic sociology, by contrast, requires a consideration of what kinds of social phenomena *could* exist and what sort of mechanisms might be required to sustain them. Building such mechanisms requires a concerted effort that spans a multiplicity of disciplines. It also, perhaps, requires a careful consideration of the political mechanisms that are responsible for choosing the sorts of social phenomena we want to bring into existence.

6 CONCLUSION

Social machines have been a prominent focus of research attention within the sciences of the Web and Internet. In this paper, we outlined a theoretical account of social machines that draws inspiration from recent work in so-called mechanical philosophy. This account—the mechanistic view—is intended to progress our conceptual understanding of what a social machine is. According to the mechanistic view, a social machine is a system whose events, states, and processes are realized by a socio-technical mechanism,

where a socio-technical mechanism is a mechanism that consists of multiple human individuals and the technological resources of the Web and Internet.

The main aim of the present paper was to explicate the nature of the relationship between social machines and CPSs. To this end, we outlined three ways of thinking about CPSs. These are what we called the compositional view, the controller view, and the constitutional view. Of these views, the constitutional view is most closely aligned with the mechanistic view of social machines. Accordingly, we suggested that a particular subset of CPSs (i.e., CPSSs) could, under the right circumstances, qualify as *bona fide* members of the class of social machines. Such social machines, we suggested, ought to be regarded as cyber-physical variants of traditional social machines. We thus drew a distinction between what we called CPSMs and traditional social machines. Some examples of CPSMs include the likes of smart cities, smart transportation systems, and environment monitoring systems. These systems are related to traditional social machines (e.g., those centered around the likes of Wikipedia and Galaxy Zoo) in the sense that their functionality is realized by socio-technical mechanisms. Nevertheless, we suggested that CPSMs accentuate some of the concerns associated with traditional social machines (e.g., issues of privacy, security, safety, and so on). They also challenge traditional approaches to system development, in the sense that they inherit many of the features (e.g., concurrency, real-time computation, and so on) that are associated with the design, development, and evaluation of CPSs.

Web science has long been recognized as an interdisciplinary endeavor. The mechanistic view does not (as far as we can tell) radically alter the kinds of disciplines that are relevant to Web science (although it does, of course, establish a particularly potent form of interdisciplinary contact with the philosophy of science). The main impact of the mechanistic view comes from the way it helps to illuminate the nature of interdisciplinary connections. In particular, the mechanistic view is nicely aligned with mechanism-based approaches in the social sciences, especially those that seek to understand social phenomena via an appeal to social mechanisms [e.g., 52]. While such approaches are typically subsumed under the general heading of analytical sociology [e.g., 24], it should be clear that social machines are, more often than not, the product of deliberate design efforts. This entails a synthetically-oriented shift in our thinking about social phenomena. Inasmuch as social machines form part of the material fabric of the social world, then they are nicely poised to serve as the building blocks for new kinds of social phenomena. In this sense, social machines are not just a lens through which we view (and hopefully understand) the nature of our social world; they are also, potentially at least, the means by which we create new social phenomena and thereby shape the structure of our social reality.

ACKNOWLEDGMENTS

This work is supported by the UK EPSRC as part of the PETRAS National Centre of Excellence for IoT Systems Cybersecurity under Grant Number EP/S035362/1 (<https://petras-iot.org/>).

REFERENCES

- [1] Dirk Ahlers, Patrick Driscoll, Erica Löfström, John Krogstie, and Annemie Wyckmans. 2016. Understanding Smart Cities as Social Machines. In *4th International*

- Workshop on the Theory and Practice of Social Machines. Montreal, Canada.
- [2] Rajeev Alur. 2015. *Principles of Cyber-Physical Systems*. MIT Press, Cambridge, Massachusetts, USA.
 - [3] Michael Baumgartner, Lorenzo Casini, and Beate Krickel. 2020. Horizontal surgicality and mechanistic constitution. *Erkenntnis* 85 (2020), 417–340.
 - [4] Tim Berners-Lee and Mark Fischetti. 1999. *Weaving the Web: The Original Design and Ultimate Destiny of the World Wide Web*. Harper Collins, New York, New York, USA.
 - [5] Christos G Cassandras. 2016. Smart cities as cyber-physical social systems. *Engineering* 2, 2 (2016), 156–158.
 - [6] Marco Conti, Andrea Passarella, and Sajal K Das. 2017. The Internet of People (IoP): A new wave in pervasive mobile computing. *Pervasive and Mobile Computing* 41 (2017), 1–27.
 - [7] Alfio Costanzo, Alberto Faro, Daniela Giordano, and Concetto Spampinato. 2016. Implementing Cyber Physical Social Systems for Smart Cities: A Semantic Web Perspective. In *13th IEEE Annual Consumer Communications & Networking Conference*. IEEE, Las Vegas, Nevada, USA, 274–275.
 - [8] Carl Craver. 2007. Constitutive explanatory relevance. *Journal of Philosophical Research* 32 (2007), 3–20.
 - [9] Carl Craver. 2007. *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience*. Clarendon Press, Oxford, UK.
 - [10] Carl Craver and James Tabery. 2016. Mechanisms in science. In *The Stanford Encyclopedia of Philosophy* (Spring 2016 ed.), Edward N Zalta (Ed.). Stanford University, Stanford, California, USA.
 - [11] Carl F Craver. 2015. Levels. In *Open MIND: Philosophy and the Mind Sciences in the 21st Century*, Thomas K Metzinger and Jennifer M Windt (Eds.). MIND Group, Frankfurt am Main, Germany, 1–26.
 - [12] Carl F Craver and Lindley Darden. 2013. *In Search of Mechanisms: Discoveries Across the Life Sciences*. The University of Chicago Press, Chicago, Illinois, USA.
 - [13] Suparna De, Yuchao Zhou, Iker Larizgoitia Abad, and Klaus Moessner. 2017. Cyber-physical-social frameworks for urban big data systems: A survey. *Applied Sciences* 7, 10 (2017), 1017.
 - [14] David De Roure, James A Hendler, Diccon James, Terhi Nurmikko-Fuller, Max Van Kleek, and Pip Willcox. 2019. Towards a Cyberphysical Web Science: A Social Machines Perspective on Pokémon GO!. In *10th ACM Conference on Web Science*, Paolo Boldi, Brooke Foucault Welles, Katharina Kinder-Kurlanda, Christo Wilson, Isabella Peters, and Wagner Jr. Meira (Eds.). ACM, Boston, Massachusetts, USA, 65–69.
 - [15] David De Roure, Clare Hooper, Megan Meredith-Lobay, Kevin Page, Ségolène Tarte, Don Cruickshank, and Catherine De Roure. 2013. Observing social machines Part 1: What to observe?. In *1st International Workshop on the Theory and Practice of Social Machines*. Rio de Janeiro, Brazil.
 - [16] Falko Dressler. 2018. Cyber physical social systems: Towards deeply integrated hybridized systems. In *International Conference on Computing, Networking and Communications*. IEEE, Maui, Hawaii, USA, 420–424.
 - [17] Xu Du, Onyeka Emebo, Aparna Varde, Niket Tandon, Sreyasi Nag Chowdhury, and Gerhard Weikum. 2016. Air quality assessment from social media and structured data: Pollutants and health impacts in urban planning. In *IEEE International Conference on Data Engineering Workshops*. IEEE, Helsinki, Finland, 54–59.
 - [18] Jon Elster. 1989. *Nuts and Bolts for the Social Sciences*. Cambridge University Press, Cambridge, UK.
 - [19] Stuart Glennan. 2017. *The New Mechanical Philosophy*. Oxford University Press, Oxford, UK.
 - [20] Stuart Glennan and Phyllis McKay Illari (Eds.). 2018. *The Routledge Handbook of Mechanisms and Mechanical Philosophy*. Routledge, New York, New York, USA.
 - [21] Bin Guo, Zhiwen Yu, and Xingshe Zhou. 2015. A data-centric framework for cyber-physical-social systems. *IT Professional* 17, 6 (2015), 4–7.
 - [22] Wendy Hall and Thanassis Tiropanis. 2012. Web Evolution and Web Science. *Computer Networks* 56 (2012), 3859–3865.
 - [23] Peter Hedström. 2005. *Dissecting the Social: On the Principles of Analytical Sociology*. Cambridge University Press, Cambridge, UK.
 - [24] Peter Hedström and Peter Bearman (Eds.). 2009. *The Oxford Handbook of Analytical Sociology*. Oxford University Press, Oxford, UK.
 - [25] Peter Hedström and Peter Bearman. 2009. What is analytical sociology all about? An introductory essay. In *The Oxford Handbook of Analytical Sociology*, Peter Hedström and Peter Bearman (Eds.). Oxford University Press, Oxford, UK.
 - [26] Peter Hedström and Petri Ylikoski. 2010. Causal mechanisms in the social sciences. *Annual Review of Sociology* 36 (2010), 49–67.
 - [27] Jim Hendler and Tim Berners-Lee. 2010. From the Semantic Web to social machines: A research challenge for AI on the World Wide Web. *Artificial Intelligence* 174 (2010), 156–161.
 - [28] James Hendler and Alice M Mulvehill. 2016. *Social Machines: The Coming Collision of Artificial Intelligence, Social Networking, and Humanity*. Apress, New York, New York, USA.
 - [29] Clare Hooper, Brian Bailey, Hugh Glaser, and James Hendler. 2016. Social machines in practice: Solutions, stakeholders and scopes. In *8th International ACM Web Science Conference*, Wolfgang Nejdl, Wendy Hall, Paolo Parigi, and Steffen Staab (Eds.). ACM, Hannover, Germany, 156–160.
 - [30] D C Horsman. 2015. Abstraction/Representation Theory for heterotic physical computing. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 373 (2015), 20140224.
 - [31] Jingwei Huang, Mamadou D Seck, and Adrian Gheorghe. 2016. Towards Trustworthy Smart Cyber-Physical-Social Systems in the Era of Internet of Things. In *11th System of Systems Engineering Conference*. IEEE, Kongsberg, Norway, 1–6.
 - [32] Phyllis McKay Illari and Jon Williamson. 2012. What is a mechanism? Thinking about mechanisms across the sciences. *European Journal for Philosophy of Science* 2, 1 (2012), 119–135.
 - [33] Edith Law and Luis von Ahn. 2011. Human Computation. *Synthesis Lectures on Artificial Intelligence and Machine Learning* 5, 3 (2011), 1–121.
 - [34] Edward A Lee. 2008. Cyber physical systems: Design challenges. In *11th IEEE International Symposium on Object and Component-Oriented Real-Time Distributed Computing*. IEEE, Orlando, Florida, USA, 363–369.
 - [35] Peter Machamer, Lindley Darden, and Carl F Craver. 2000. Thinking about mechanisms. *Philosophy of Science* 67, 1 (2000), 1–25.
 - [36] Kieron O'Hara. 2012. Trust in social machines: The challenges. In *AISB/IACAP World Congress 2012: Social Computing, Social Cognition, Social Networks and Multiagent Systems*. Birmingham, UK.
 - [37] Spyridon Orestis Palermos. 2017. Social Machines: A Philosophical Engineering. *Phenomenology and the Cognitive Sciences* 16, 5 (2017), 953–978.
 - [38] Gualtiero Piccinini. 2015. Computation in physical systems. In *The Stanford Encyclopedia of Philosophy* (Summer 2015 ed.), Edward N Zalta (Ed.). Stanford University, Stanford, California, USA.
 - [39] Raj Rajkumar, Dionisio de Niz, and Mark Klein (Eds.). 2017. *Cyber-Physical Systems*. Addison-Wesley, Boston, Massachusetts, USA.
 - [40] Ragunathan Rajkumar, Insup Lee, Lui Sha, and John Stankovic. 2010. Cyber-physical systems: The next computing revolution. In *Design Automation Conference*, Sachin Sapatnekar (Ed.). ACM, Anaheim, California, USA, 731–736.
 - [41] Nigel R Shadbolt. 2013. Knowledge acquisition and the rise of social machines. *International Journal of Human-Computer Studies* 71, 2 (2013), 200–205.
 - [42] Nigel R Shadbolt, Kieron O'Hara, David De Roure, and Wendy Hall. 2019. *The Theory and Practice of Social Machines*. Springer, Cham, Switzerland.
 - [43] 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 *1st International Workshop on the Theory and Practice of Social Machines*. Rio de Janeiro, Brazil.
 - [44] Paul R Smart, Aastha Madaan, and Wendy Hall. 2019. Where the Smart Things Are: Social Machines and the Internet of Things. *Phenomenology and the Cognitive Sciences* 18, 3 (2019), 551–575.
 - [45] Ségolène Tarte, David De Roure, and Pip Willcox. 2014. Working out the plot: The role of Stories in Social Machines. In *2nd International Workshop on the Theory and Practice of Social Machines*. Seoul, Korea.
 - [46] Ramine Tinati and Leslie Carr. 2012. Understanding social machines. In *ASE/IEEE International Conference on Social Computing and International Conference on Privacy, Security, Risk and Trust*. Amsterdam, The Netherlands.
 - [47] Oksana Tokarchuk, Roberta Cuel, and Marco Zamarian. 2012. Analyzing Crowd Labor and Designing Incentives for Humans in the Loop. *IEEE Internet Computing* 16, 5 (2012), 45–51.
 - [48] Francisco J Valero-Cuevas, Jae-Woong Yi, Daniel Brown, Robert V McNamara, Chandana Paul, and Hood Lipson. 2007. The tendon network of the fingers performs anatomical computation at a macroscopic scale. *IEEE Transactions on Biomedical Engineering* 54, 6 (2007), 1161–1166.
 - [49] Fei-Yue Wang. 2010. The emergence of intelligent enterprises: From CPS to CPSS. *IEEE Intelligent Systems* 25, 4 (2010), 85–88.
 - [50] Robert A Wilson and Carl Craver. 2007. Realization: Metaphysical and Scientific Perspectives. In *Philosophy of Psychology and Cognitive Science*, Paul Thagard (Ed.). North-Holland, Oxford, UK, 81–104.
 - [51] Gang Xiong, Fenghua Zhu, Xiwei Liu, Xisong Dong, Wuling Huang, Songhang Chen, and Kai Zhao. 2015. Cyber-physical-social system in intelligent transportation. *IEEE/CAA Journal of Automatica Sinica* 2, 3 (2015), 320–333.
 - [52] Petri Ylikoski. 2018. Social mechanisms. In *The Routledge Handbook of Mechanisms and Mechanical Philosophy*, Stuart Glennan and Phyllis McKay Illari (Eds.). Routledge, New York, New York, USA, 401–412.