

Applying Mechanical Philosophy to Web Science: The Case of Social Machines

Paul R. Smart · Kieron O'Hara · Wendy
Hall

Received: date / Accepted: date

Abstract Social machines are a prominent focus of attention for those who work in the field of Web and Internet science. Although a number of online systems have been described as social machines (examples include the likes of Facebook, Twitter, Wikipedia, Reddit, and Galaxy Zoo), there is, as yet, little consensus as to the precise meaning of the term “social machine.” This presents a problem for the scientific study of social machines, especially when it comes to the provision of a theoretical framework that directs, informs, and explicates the scientific and engineering activities of the social machine community. The present paper outlines an approach to understanding social machines that draws on recent work in the philosophy of science, especially work in so-called mechanical philosophy. This is what might be called a mechanistic view of social machines. According to this view, social machines are systems whose phenomena (i.e., events, states, and processes) are explained via an appeal to (online) socio-technical mechanisms. We show how this account is able to accommodate a number of existing attempts to define the social machine concept, thereby yielding an important opportunity for theoretical integration.

Keywords Social Machine · Web Science · Mechanism · Socio-Technical System · Mechanistic Explanation · Internet · Philosophy of Web Science

Paul R. Smart
Electronics & Computer Science, University of Southampton, Southampton, SO17 1BJ, UK.
E-mail: ps02v@ecs.soton.ac.uk

Kieron O'Hara
Electronics & Computer Science, University of Southampton, Southampton, SO17 1BJ, UK.
E-mail: kmo@ecs.soton.ac.uk

Wendy Hall
Electronics & Computer Science, University of Southampton, Southampton, SO17 1BJ, UK.
E-mail: wh@ecs.soton.ac.uk

1 Introduction

A key trend in the recent technological evolution of the World Wide Web (henceforth the Web) has been the development of applications and services that support greater levels of user participation in the generation and management of online content. From its origins as a platform for information dissemination, the Web has evolved into a highly interactive environment—one that supports a wide range of information processing tasks and socially-oriented activities. This has given rise to a rich array of phenomena that are of interest to a number of scientific disciplines. These phenomena are typically not ones that can be understood by considering the purely social or technological aspects of the Internet and Web; rather, they require an interdisciplinary approach, an approach that is epitomized by the nascent discipline of Web and Internet Science (WAIS) (Berners-Lee et al. 2006a,b; Tiropanis et al. 2015).

One of the topics that has been at the forefront of recent WAIS research relates to what are called *social machines* (Hendler and Berners-Lee 2010; Hendler and Mulvehill 2016; Shadbolt et al. 2019). These are typically viewed as online (i.e., Web- or Internet-based) systems that support large-scale forms of social participation. Some of the systems that have been the target of social machine research include the likes of Facebook (Burégio et al. 2013b; Hendler and Berners-Lee 2010), Twitter (Hall and Tiropanis 2012), YouTube (Shadbolt et al. 2013), Ushahidi (Hendler and Mulvehill 2016; Shadbolt 2013), Galaxy Zoo (Hall and Tiropanis 2012; Shadbolt 2013), reCAPTCHA (O’Hara 2012), Wikipedia (Hall and Tiropanis 2012; Hendler and Berners-Lee 2010; Shadbolt 2013), Pokémon GO! (De Roure et al. 2019), Reddit (Hooper et al. 2016; Shadbolt et al. 2013), Flickr (Shadbolt et al. 2013), and the Web itself (Hall and Tiropanis 2012). As should be clear from this list, social machines occupy a variety of functionally-diverse niches within the ecology of the Web and Internet. Extant social machines thus include social networking systems (Facebook, Twitter), microblogging services (Twitter), video/photo sharing systems (YouTube, Flickr), citizen science projects (Galaxy Zoo), social news sites (Reddit), collaborative content editing sites (Wikipedia), frameworks for the creation of collaborative systems (Ushahidi), augmented reality games (Pokémon GO!), and systems that support the productive exploitation of human contributions in the context of conventional computational processes (reCAPTCHA).

Social machines are clearly an important focus of empirical and theoretical attention for the WAIS community. There is, however, a pressing problem that lies at the heart of the social machine research effort. This problem relates to the precise meaning of the term “social machine.” There is, in fact, little consensus at the present time as to how this term should be understood. This does not mean that there is no consensus regarding the general properties of social machines. Social machines are thus typically presented as Web- or Internet-based systems that work to support social interaction and collaboration (often at large scale). But beyond this general notion of (large-scale) technology-mediated social participation, there is little agreement as to what it is that

makes something a social machine. This is important, because in the absence of an ability to say what social machines are it is difficult to identify the precise empirical targets of the social machine research effort.

In this paper, we outline an approach to understanding social machines that goes by the name of the *mechanistic view*. This approach draws inspiration from work in the philosophy of science, especially work that goes under the heading of mechanical or neo-mechanical philosophy (Glennan 2017; Glennan and Illari 2018). In essence, we attempt to provide a philosophical account of the social machine concept by drawing on analytical resources that were originally developed for theories of mechanistic explanation in the philosophy of science. This is what might be called a neo-mechanical approach to social machines. The present paper highlights some of the advantages of this approach. In Section 2, we provide an overview of existing efforts to define the social machine concept. Section 3 then describes the key features of a neo-mechanical approach to social machines, an approach we label the mechanistic view. In subsequent sections, we show how the mechanistic view is able to accommodate some existing accounts of social machines, thereby providing the basis for theoretical integration. We also intimate at some of the advantages afforded by the mechanistic view. In particular, we show how the mechanistic view has the potential to inform approaches to social machine taxonomization (see Section 4), highlight interdisciplinary linkages (see Section 4 and Section 8), advance our understanding of some of the intuitions that have been expressed by the social machine community (see Section 6), and shed light on some of the epistemic practices of the social machine community (see Section 8).

The primary goal of the present paper is to show how recent work in the philosophy of science can be used to support our understanding of the social machine concept. At a more general level, however, we aim to draw attention to the importance of socio-technical mechanisms as a focus area for future work in the philosophy of (Web) science.

2 Social Machines

Social machines were first introduced into the WAIS literature by Berners-Lee and Fischetti (1999). 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. (Berners-Lee and Fischetti 1999, p. 172, original emphasis)

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 addition to the content creation view, a number of other views of social machines have emerged in recent years. Table 1 provides a summary of all the views that we will consider in the present paper. As can be seen from this table, there are eight views in total, including the aforementioned content creation view. For the purposes of discussion, we will organize these views into a number of categories (see Table 1). *Activity views*, for example, seek to impose constraints on the activities of the constituents of a social machine. The content creation view is, as far as we can tell, the only member of this category at the present time.

A second category of views are what might be called *functional views*. Members of this category tend to focus their attention on the kinds of things a social machine does. Proponents of functional views generally accept that social machines are, at root, socio-technical systems; however, they seek to limit the sort of systems that can qualify as social machines by imposing constraints on the processes performed by these systems. 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 regard a social machine as a form of distributed cognitive system, i.e., as a system in which cognitive processing routines (e.g., problem-solving processes) are distributed across multiple human agents and technological artifacts. Social machines have also been approached from a socio-computational perspective (Horsman 2015; Luczak-Roesch et al. 2015; Strohmaier and Wagner 2014). Horsman (2015), for example, characterizes social machines as “computational ecosystems comprising both digital computers and multiple human users, all acting towards a computational goal” (Horsman 2015, p. 2). Other proponents of the socio-computational view include Strohmaier and Wagner (2014). They suggest that social machines are a species of socio-computational system “whose emergent properties are co-determined by the social behavior of their users and by [the] algorithmic computation of machines” (p. 84).

A third category of views are what we have labelled *bare occurrent views*. These views are somewhat unlike other views in the sense that they characterize social machines as either tasks or processes. This contrasts with the way that social machines are typically referred to as physical systems in the context of other views. From a metaphysical standpoint, the difference between bare occurrent views and other views relates to the distinction between occurrents and continuants (see Kaiser and Krickel 2017). Occurrents are things like processes, states, and events. Continuants, by contrast, are objects like chairs, stones, tables, and so on. While many of the existing views of social machines appear to acknowledge the status of social machines as physical systems or objects (that is to say, continuants), proponents of bare occurrent views (e.g., the process view) tend to regard social machines as something more akin to occurrents.

Table 1 Existing views of social machines.

View	Category	Description	Source
Content Creation View	Activity	A social machine is a Web-based system in which people do the creative work and the technology does the administration.	Berners-Lee and Fischetti (1999)
Problem-Solving View	Functional	A social machine is a problem-solving system.	Hooper et al. (2016)
Cognitive Systems View	Functional	A social machine is a form of distributed cognitive system in which cognitive processing routines are distributed across multiple human individuals and a suite of (Web-based) technological components.	Palermos (2017)
Socio-Computational View	Functional	A social machine is a socio-computational system in which some aspects of a computational process or routine are delegated to multiple human individuals.	Luczak-Roesch et al. (2016)
Process View	Bare Occurrent	A social machine is a task or process that involves the participation of multiple human individuals.	Tinati and Carr (2012)
Sociable Machines View	Socialized Technology	A social machine is a technological system that is able to interact with other technological systems within the larger socio-technical environment of the Web or Internet.	Burégio et al. (2013a)
Socially-Situated Machines View	Socialized Technology	Social machines are intelligent (technological) systems whose processing capabilities are grounded in their exploitation of social information or their interactions with human agents.	Smart and Madaan (2017)
Narratological View	Methodological	Social machines are systems that are amenable to narratological analysis.	Tarte et al. (2014)

The sociable machines view and the socially-situated machines view are both examples of what we call *socialized technology views*. Views of this stripe tend to regard social machines as purely technological systems, albeit ones that possess certain social properties or capabilities. Meira et al. (2011) are among the most vocal champions of the sociable machines view. They suggest that we ought to conceive of social machines as a form of “sociable software” or “social service component”—a particular kind of application that is able to interoperate with other applications and systems. According to Meira et al. (2011),

this emphasis on inter-operability and the capacity to form ‘relationships’ with other systems makes sociability of crucial importance to our understanding of social machines:

By the very nature of the concept we are proposing, SMs [social machines] are sociable stuff and, in nearly all cases, each one should provide [a] means to interact with one another. The concept of relationships between SMs is similar to that of relationships between people; we can view them as trusted relations between different SMs, satisfying established constraints, which describe any restrictions that a given SM can have. (Meira et al. 2011, p. 2)

As this quote makes clear, the relationships that exist between social machines are deemed to resemble those encountered in the context of conventional human–human relationships (see also Burégio et al. 2013b). In particular, notions such as trust, reciprocity, reliability, intent, reputation, credibility, accountability, compatibility, sociability, honesty, responsivity, and so on, all seem to be important in describing the characteristics of social machines.

The socially-situated machines view resembles the sociable machines view, in the sense that it regards social machines as technological systems. In contrast to the sociable machines view, however, the socially-situated machines view tends to direct more attention to the way in which the (cognitive) capacities of a technological system are acquired. Perhaps the most explicit expression of the socially-situated machines view is provided by Smart and Madaan (2017). They suggest that the Web provides an unprecedented form of contact with the human social environment, one that enables machine learning systems to capitalize on various forms of socially-supported learning and socially-scaffolded development. The Internet and Web, Smart and Madaan (2017) suggest, extend the reach of socially-scaffolded development to the realm of machine intelligence, thereby yielding synthetic cognitive systems that are as much a product of social forces and factors as are the human minds that created them.

The final category of views are methodological views. Members of this category are not so much an attempt to state what social machines are as they are an attempt to outline the virtues of a particular methodological approach to understanding social machines. The narratological view, which we equate with the work of Tarte et al. (2014), is, at present, the only member of this category. Tarte et al. (2014) advocate a “storytelling approach” to social machines. They suggest that existing work has tended to overlook the social aspects of social machines and that a storytelling approach can help to address this shortcoming.

We have now completed our survey of the different ways that social machines have been conceptualized in the relevant scientific literature. At first sight, these views are a highly mixed bunch and it is far from clear that they can all be accommodated under a common conceptual umbrella. On a more positive note, however, it should be noted that all the views discussed above share some common features. There is, for example, a degree of consensus regarding the importance of *active* social participation. Systems in which a collection of human agents merely browse or consume content without contributing anything

in return do not seem to be regarded as genuine members of the class of social machines. This suggests that a unified theoretical account of social machines will need to respect the idea that people are actively involved in some process, task, or activity. This is what might be called the *social participation constraint*.

Another constraint relates to the *online* nature of social machines. For the most part, social machines are conceived as Web- or Internet-based systems, which is to say they are conceived as online systems. As such, a theoretical account will need to accommodate the idea that social machines are some form of online system. Let us call this the *online system constraint*.

In addition to these two (strict or hard) constraints, a theoretical account of social machines might be expected to possess some other features. These are perhaps better glossed as ‘optional extras’ (or soft constraints) as opposed to hard constraints; nevertheless, it will be helpful to cast them as hard constraints for the purpose of evaluating attempts at theoretical unification. The constraint we will consider here is what is called the *backwards compatibility constraint*. The *backwards compatibility constraint* mandates that a new account of the social machine concept ought to be consistent with the content creation view, in the sense that all social machines identified via the content creation view will continue to be recognized as social machines by the new account (this is what we mean by backwards compatibility). This is important because the content creation view is the most popular view of social machines, and it has played an important role in guiding research efforts for the better part of decade. With this in mind, it seems important that any new theoretical account of social machines ought to accommodate the body of existing work that has been undertaken under the rubric of the content creation view.

Having identified some of the constraints that ought to be respected by a theoretical account of social machines, let us now turn our attention to the details of an account that speaks to such constraints. In the next section, we will present a new view of social machines—the *mechanistic view*. This view is intended to satisfy the three constraints discussed above (i.e., the social participation constraint, the online system constraint, and the backwards compatibility constraint). It also, we suggest, accommodates many of the views presented in Table 1, thereby providing us with an important opportunity for theoretical integration.

3 The Mechanistic View

According to the mechanistic view, social machines are to be understood as online systems whose phenomena (e.g., events, states, processes, behaviors, and activities) are realized by socio-technical mechanisms, where the notion of a socio-technical mechanism is to be understood as a mechanism comprising 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* or *neo-mechanical philosophy* (Glennan 2017; Glennan and Illari 2018). 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 (Craver and Tabery 2016). 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,” there appears to be a general consensus about the basic features of mechanisms. This consensus is captured in the following succinct definition of the mechanism concept:

A mechanism for a phenomenon consists of entities and activities organized in such a way that they are responsible for the phenomenon. (Illari and Williamson 2012, p. 120)

As is clear from this definition, the building blocks of mechanisms are what are called “entities” and “activities.” These building blocks are what are sometimes referred to as the components of mechanisms. While there has been some debate regarding the extent to which entities and activities ought to be seen as distinct components, the position we will adopt here regards the term “component” as referring to both entities and activities. That is to say, we regard a component as an entity involved in an activity, as opposed to an alternative view, which views entities and activities as two distinct kinds of component.¹ 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” (Machamer et al. 2000, p. 3) and as the “causal components of mechanisms” (Craver 2007b, p. 6). They “are the things that the entities do” (Craver and Darden 2013, 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).

In the context of the mechanistic view, we are interested in mechanisms of a particular type, namely, socio-technical mechanisms. Such mechanisms are individuated by the nature of the entities that comprise the mechanism. In particular, a socio-technical mechanism is a mechanism whose entities are drawn from both the social and the technological realms. These are what we will call *social components* and *technological components*.² For present purposes,

¹ This alternative view is commonly referred to as the *duality thesis* in mechanical philosophy (Kaiser 2018). See Glennan (2017, p. 21), for further discussion of this issue.

² As noted above, the term “component” refers to both entities and activities. In the present case, we are assuming that the status of a component as a specifically social or

we can think of social components as human individuals and technological components as digital computational systems (or devices). Accordingly, we can conceive of a socio-technical mechanism as a mechanism in which human agents (*qua* social components) and digital computers (*qua* technological components) interact in such a way as to be responsible for some kind of phenomenon.

At this point, we need to address two issues. The first relates to the notion of *socio-technical integration*. In particular, what is it that determines the hybrid (i.e., socio-technical) status of a mechanism? What is it, in other words, that makes social and technological components part of a common, integrated mechanism?

The answer to this question requires us to recognize a distinction between two forms of mechanistic explanation. Following the seminal work of Salmon (1984), many philosophers recognize a distinction between etiological (or causal) mechanistic explanation and constitutive mechanistic explanation (Craver 2007b; Kaiser 2018; Kaiser and Krickel 2017; Ylikoski 2013). The difference between these two kinds of mechanistic explanation is that, in the etiological case, the mechanism as a whole is said to cause (or produce) the phenomenon of interest. In the constitutive case, by contrast, the mechanism is said to constitute or realize the phenomenon of interest.³ The difference between etiological and constitutive mechanistic explanation thus turns on the nature of the relation that exists between mechanisms and the phenomena they explain. In the etiological case, this relation tracks a causal dependency between the mechanism and the phenomenon, and we thus understand why the phenomenon occurred with respect to the mechanism that caused or produced it. A mechanism is thus deemed to be *responsible* for a phenomenon (recall the above definition of the mechanism concept) in the sense that the mechanism caused the phenomenon to occur. This contrasts with the case of constitutive mechanistic explanation, where the aim is to understand a phenomenon by describing the mechanisms that underlie, realize, or constitute the phenomenon. In the constitutive case, then, a mechanism is deemed to be *responsible* for a phenomenon in the sense that it constitutes or realizes the phenomenon.

This distinction between etiological and constitutive mechanistic explanation is paralleled by a corresponding distinction between two forms of explanatory relevance, namely, *causal relevance* and *constitutive relevance* (see Craver 2007a). Causal relevance identifies the material objects that help to cause, produce, or bring about a phenomenon. Accordingly, the notion of causal relevance is of crucial importance when it comes to etiological mechanistic explanations. Constitutive relevance, by contrast, identifies the material objects

technological component is determined by the nature of the *entity* that performs an activity. A social component is thus a human individual involved in some activity, and a technological component is a computational system involved in some activity.

³ Note that in the present paper, we will assume, without further clarification, that the notion of *mechanistic constitution* (e.g., Baumgartner et al. 2020) is semantically-equivalent to the notion of *mechanistic realization* (e.g., Wilson and Craver 2007). See Harinen (2018), for more on this.

that are constitutively relevant to a phenomenon, and it is thus important in the case of constitutive mechanistic explanations.

From the standpoint of the mechanistic view, issues of socio-technical integration are resolved via an appeal to issues of constitutive relevance. It is thus not enough to say that a particular phenomenon was caused by the activity of human individuals and computational systems. Instead, what we are looking for is a state-of-affairs in which human individuals and computational systems form part of a common physical fabric that realizes or constitutes a particular phenomenon. In other words, what makes human individuals and computational systems part of a common mechanism is that they are constitutively relevant to a single explanandum phenomenon. What this means, in effect, is that the social (human) and technological (computational system) components of a mechanism 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 exhibited by a conventional automobile.

This account of socio-technical integration speaks to some of the intuitions expressed by members of the social machine community. Within the social machine literature, we thus encounter the idea of social machines as “integrated systems of people and computers” (Strohmaier 2013), as “purposefully designed sociotechnical system[s] comprising machines and people” (De Roure et al. 2013), as systems in which “the human and digital parts...[form] a machine in which the two aspects are seamlessly interwoven” (Shadbolt et al. 2013), as “large-scale sociotechnical systems that combine human and machine computations and interactions into a new system with emergent characteristics” (Ahlers et al. 2016), as systems in which the “components...are people and technologies” (O’Hara 2012), as systems that involve “the co-constitutional involvement of humans and technologies” (Tinati and Carr 2012), as systems whose “behaviour [is] co-created by human participants and technological components” (De Roure et al. 2019), and as systems that are constituted by “humans, machines, and their interactions” (Tarte et al. 2014). Courtesy of the appeal to socio-technical mechanisms (and constitutive relevance), the mechanistic view is able to accommodate all these intuitions.

In addition to issues of socio-technical integration, it is important to say something about the kinds of phenomena that a socio-technical mechanism is deemed to be responsible for. In other words, what kinds of phenomena are realized by socio-technical mechanisms under the mechanistic view?

Thankfully, the answer, in this case, is straightforward: The mechanistic view is utterly neutral as regards the kinds of phenomena that are realized by socio-technical mechanisms. In a social machine context, the phenomena of interest are likely to be those of the cognitive, social, or computational kind. These distinctions, however, are unimportant from the standpoint of the mechanistic view; at the very least, such distinctions play no role in determining what kinds of systems ought to be counted as social machines. In general, the targets of (constitutive) mechanistic explanations (across all the sciences) are what Kaiser and Krickel (2017) refer to as *object-involving-occurrents*,




Fig. 1 A graphical depiction of a phenomenon (top) and its mechanism (bottom) (adapted from Craver 2007b, p. 7). (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. This mechanism consists of multiple human individuals and one or more technological systems/devices.

which refers to the events, states, and processes (the occurrents) associated with a particular physical object, entity, or system (the object). According to Kaiser and Krickel (2017), object-involving-occurrents consist of “an object (or system) that is engaged in a certain occurrent” (p. 768). In the context of the mechanistic view, the object (or system) denotes a social machine, but no constraints are placed on the kinds of occurrents (i.e., the kinds of events, states, and processes) in which a social machine is involved.

To help us understand the properties of socio-technical mechanisms from the standpoint of the mechanistic view, it will be useful to refer to a diagrammatic scheme that is commonly attributed to Craver (2007b) (these are sometimes referred to as *Craver diagrams*). Figure 1a depicts a generic mechanism, which is represented according to this scheme. Craver refers to the mechanism as S and the constituents 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 the components of the mechanism. The directed arrows represent the causal interactions between the components, and it is these causal interactions that are deemed to be responsible (in a constitutive sense) for S ’s Ψ -ing.⁴

⁴ As noted by an anonymous reviewer, it is unclear whether the arrows in Craver’s (2007b) diagram are intended to represent causal relationships. While Craver (2007b, p. 7) suggests that the arrows represent activities, a number of philosophers have interpreted the arrows as representing causal relationships. These include Harinen (2018, p. 40) (“The relations within the levels, represented by arrows, are causal.”), Clarke et al. (2014, p. 1654) (“Solid arrows represent intra-level causal relations.”), and Baumgartner and Gebharder (2016, p. 734) (“The arrows symbolize direct causation. . .”). This particular point of confusion regarding Craver’s diagram is referred to by Menzies (2012). In particular, he notes that: “Craver says that the activities are represented by the Φ -ings. . . and in the same paragraph that they are represented by the arrows. He must mean that the Φ -ings are the activities, the arrows

Figure 1b depicts a socio-technical mechanism using the same diagrammatic scheme as that used in Figure 1a. 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 (or, more generally, social agents [see Section 4]). This is what we will refer to as the *sociality criterion*. The sociality criterion is intended to satisfy the social participation constraint discussed in Section 2. In particular, the sociality criterion ensures that multiple human individuals will be included as social components in a socio-technical mechanism. In addition, the sociality criterion ensures that these components must be active, in the sense that they engage in causal interactions with at least one other component in the *same* mechanism. If a component does not interact with at least one other component of the same mechanism, then it cannot be constitutively relevant to the phenomenon that is realized by that mechanism. This is, in fact, what active social participation means from the standpoint of the mechanistic view: it is a claim about the effect that a human agent’s activity has on other components of the same mechanism. The upshot is that the social participation criterion is implicitly satisfied by the idea that a socio-technical mechanism must consist of *multiple* human agents, each of which qualifies as a component of the mechanism.

The arrows in Figure 1b represent the causal links between the components of a socio-technical mechanism. 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 these interactions are permissible under the mechanistic view. The only constraint imposed on the causal relationships is that *some* of these relationships must involve the use of Internet protocols. This is what we will dub 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 systems (or, more informally, online systems). The IP criterion is intended to satisfy the online system constraint that was discussed in Section 2.⁵

are causal relations, and that the activities logically imply the existence of causal relations” (Menzies 2012, p. 797, footnote 1).

⁵ Note that the IP criterion applies to systems that qualify as both Web- and Internet-based systems. 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 the HyperText Transfer Protocol (HTTP). Examples of such systems include Facebook, Wikipedia, Twitter, Galaxy Zoo, and so on. In addition to these Web-based systems, it has been suggested that the social machine concept is also applicable to Internet-based systems (Ahlers et al. 2016; De Roure et al. 2013; Hooper et al. 2016; Smart et al. 2019). 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). The IP criterion is sufficiently

To summarize the mechanistic view: Social machines are systems whose phenomena (events, states, or processes) are realized by socio-technical mechanisms, and the notion of a socio-technical mechanism is to be understood as a particular kind of (hybrid) mechanism that is composed of entities drawn from both the social (e.g., human individuals) and technological (e.g., conventional digital computers) realms. Accordingly, from the standpoint of the mechanistic view, socio-technical mechanisms are of crucial importance in determining whether or not some online system ought to be counted as a *bona fide* social machine. Such mechanisms are typically revealed as part of the attempt to provide (mechanistically-oriented) explanations of online phenomena, i.e., the sort of phenomena that fall within the intellectual and empirical remit of the WAIS community. It should be noted, however, that WAIS is, in part, a synthetic enterprise—an attempt to build (or synthesize) online systems as opposed to merely analyzing existing systems. Thus, while the mechanistic view of social machines is inspired by theories of mechanistic explanation within the philosophy of science, this does not mean that social machines are *only* revealed as the result of attempts at mechanism discovery; they can also emerge as part of the effort to implement socio-technical mechanisms and thereby instantiate online phenomena.

At this point, it should be clear that the mechanistic view is not a particularly straightforward account of social machines—it requires an understanding of concepts that will no doubt be unfamiliar to many of those who study (and/or build) Web- or Internet-based systems. Having said this, it should be noted that the foregoing discussion is nothing more than a philosophical characterization of the explanatory practices of scientists across a broad array of scientific (and engineering) disciplines (see Glennan and Illari 2018). In attempting to explain a particular phenomenon, scientists often attempt to identify the mechanisms that are responsible for the phenomenon in question (Craver and Darden 2013). In explaining cognitive phenomena (e.g., spatial memory), for example, a cognitive scientist will attempt to detail the actual mechanisms that are responsible for the phenomena (e.g., the mechanisms that mediate changes in synaptic plasticity). In the case of social machines, the mechanisms of interest are socio-technical mechanisms, i.e., mechanisms that consist of both social and technological components. But apart from this crucial difference, there is no reason to see the mechanistic view as committing social machine researchers to explanatory practices that differ from those of any other scientific discipline. As with other areas of scientific inquiry, the effort to understand social machines is a research effort that is concerned (at least in part) with the discovery, identification, and analysis of mechanisms. This is the case, regardless of whether or not we accept the mechanistic view. All that the mechanistic view requires us to do is direct our attention to mechanisms of a particular kind. These mechanisms are what we have called socio-technical mechanisms.

generic to accommodate research into both Web-based systems and systems that are built around the use of IoT devices.

4 Conceptual Liberalism

As we have seen, socio-technical mechanisms are a key feature of the mechanistic view, and they limit the kinds of systems that gain admittance to the social machine club. That being said, the mechanistic view should not be regarded as an overly stringent gatekeeper. Aside from the appeal to socio-technical mechanisms, there is nothing to prevent a rich diversity of systems being the focus of scientific attention for the WAIS community. This stems from the way in which the notion of a socio-technical mechanism is defined under the mechanistic view. Such mechanisms, recall, are simply mechanisms that consist of two or more social components and one or more technological components, with IP-compatible networks helping to support some of the causal interactions between the components. This way of defining a socio-technical mechanism yields a variety of forms of liberalism. The first form of liberalism is what we will call *activity liberalism*. Activity liberalism refers to the activities performed by the entities of a mechanism. Under the mechanistic view, no constraints are placed on the activities of the entities of a socio-technical mechanism: social and technological entities are free to engage in any kind of activity. The second form of liberalism is what we will call *phenomenal liberalism*. This refers to the nature of the phenomenon that is realized by a socio-technical mechanism. Under the mechanistic view, socio-technical mechanisms can realize phenomena of any kind (e.g., social phenomena, computational phenomena, epistemic phenomena, cognitive phenomena, political phenomena, economic phenomena, and so on). Finally, the mechanistic view is somewhat liberal regarding the nature of social components. Under the mechanistic view, social components are conceived as human individuals (involved in some activity), but it is perhaps more appropriate to refer to the social components as *social agents*. Whether the notion of a social agent ought to be seen as synonymous with the notion of a human agent is something that is deserving of further scrutiny. For present purposes, however, we will assume that the social components of socio-technical mechanisms are social agents and that the notion of a social agent is somewhat broader than the notion of a human agent. In short, we allow for the possibility that non-human agents can qualify as the social components of socio-technical mechanisms. This is what we will refer to as *social liberalism*.⁶

Each of these kinds of liberalism are important, for they enable us to see how the mechanistic view might serve as an overarching theoretical framework that accommodates existing views of social machines. In respect of activity liberalism, for example, the components of a socio-technical mechanism are permitted to engage in the sorts of activities that are the mainstay of the

⁶ Some will no doubt be inclined to see this as a weakness of the mechanistic view. In particular, the failure to provide a robust account of what it means to be a social agent may be seen as a shortcoming of the mechanistic view. This is, to be sure, an important issue, but (for reasons of space) it is one that we must defer to future work. For present purposes, we can think of human agents as paradigmatic examples of social agents, but it is unclear whether the class of social agents ought to be deemed as coextensive with the class of human agents. A social robot, for example, may qualify as a *bona fide* social agent despite the fact that it fails to meet the criteria for humanhood (or, perhaps, personhood).

content creation view. The content creation view, recall, distinguishes between the activities of the human and technological elements of a social machine, with the humans engaged in creative work and the technological components engaged in administrative work. From the standpoint of the mechanistic view, there is nothing to prevent systems of this type from qualifying as social machines. The mechanistic view is simply agnostic with regard to the activities (or functional roles) performed by the various components of a socio-technical mechanism; it does not rule out the possibility that in some cases (and perhaps a great many cases) the functional roles of the social and technological components will comply with the sorts of activity-related constraints imposed by the content creation view. Accordingly, there is no reason to regard the mechanistic view as incompatible with the content creation view. Evidently, the two views are not the same: the mechanistic view embraces activity liberalism, whereas the content creation view does not. But apart from this important difference, there is no reason to regard the content creation view as radically distinct from the mechanistic view. The content creation view merely draws attention to the distinctive roles that might be performed by the components of a socio-technical mechanism in a particular application context. In this sense, the mechanistic view might be regarded as a generalization of the content creation view. The mechanistic view thus respects the underlying intuition that social machines are to be understood as (online) socio-technical systems, but it rejects the call for activity-related constraints to be applied to the constituents of a socio-technical mechanism. In this sense, the mechanistic view satisfies the backwards compatibility constraint discussed in Section 2: it allows systems that were identified as social machines by the content creation view to still count as *bona fide* members of the social machine club.

Let us now direct our attention to the second form of liberalism introduced above, namely, phenomenal liberalism. The advantage of phenomenal liberalism is that it frees us from concerns about the extent to which social machines ought to be individuated with respect to the nature of the phenomena (e.g., processes) in which they are involved. Consider, for example, the idea that social machines ought to be regarded as (distributed) cognitive systems (Palermos 2017) or socio-computational systems (Horsman 2015). Inasmuch as these kinds of distinctions are made with respect to the nature of the phenomena exhibited by a system (e.g., a cognitive system is to be understood as a system that exhibits cognitive phenomena), then they are all permissible under the mechanistic view. The mechanistic view is thus sufficiently generic to accommodate the idea that social machines might exist as (e.g.) socio-computational systems (Horsman 2015), (distributed) cognitive systems (Palermos 2017), knowledge machines (Smart 2018), and mathematical social machines (Martin 2016). There is, in addition, nothing to prevent a social machine from functioning as multiple kinds of systems at the same time. A social machine could thus function as a mathematical social machine and as a cognitive system on the grounds that (e.g.) some episode of mathematical problem-solving qualifies as a form of socially-distributed cognitive processing. Similarly, inasmuch as the relevant process leads to new insights and discoveries, and thus contributes

to our growing body of mathematical knowledge, then it is perfectly possible for a social machine to count as both a mathematical social machine and as a knowledge machine.

At this point, some of the benefits of phenomenal liberalism start to come into sharper focus. One of these benefits relates to the potential for theoretical integration. In particular, phenomenal liberalism allows us to accommodate both the cognitive systems view and the socio-computational view within the theoretical framework of the mechanistic view. The mechanistic view thus allows for the idea that socio-technical mechanisms may, on occasion, serve as the realization base for phenomena that qualify as either cognitive or computational in nature. There is, as such, no reason to see the cognitive systems view or the socio-computational view as in any sort of tension with the mechanistic view.

A second benefit of phenomenal liberalism relates to the effort to taxonomize social machines. The existing literature refers to social machines of various types. These include mathematical social machines (Martin 2016), knowledge machines (Smart 2018), and musical social machines (Shadbolt et al. 2019, pp. 143–144). From the standpoint of the mechanistic view, these distinctions can be understood as an attempt to discriminate categories of social machines according to the nature of the *phenomenon* that is realized by a socio-technical mechanism or, alternatively, the thing (e.g., the artifact) that is produced by a socio-technical mechanism. A mathematical social machine might thus be conceived as a social machine that engages in some form of mathematical activity, while a musical social machine might be conceived as a social machine that (courtesy of its operation) produces a musical artifact.

A third benefit of phenomenal liberalism relates to the interdisciplinary character of the social machine research effort. Because the mechanistic view places no constraints on the nature of the phenomena that are realized by socio-technical mechanisms, it is able to accommodate the idea that social machines are just as relevant to those who study *social* phenomena as they are to those who study *cognitive* or *computational* phenomena. What is more, such forms of interdisciplinary contact are revealed precisely as a result of the theoretical character of the mechanistic view. The linkages with other disciplines thus emerge not just because we are able to characterize phenomena in cognitive, computational, or social terms; rather, they emerge as a direct result of the commitment to mechanistic explanation and mechanism-driven modes of scientific discovery (see Craver and Darden 2013) across a multiplicity of different disciplines, including sociology (Hedström 2005; Hedström and Ylikoski 2010; Ylikoski 2018), computer science (especially, cyber-security) (Hatleback and Spring 2014; Spring and Hatleback 2017; Spring and Illari 2019; Spring et al. 2017), cognitive science (Bechtel 2008; Zednik 2018), and engineering (van Eck 2018). Relative to this body of existing work, the mechanistic view is revealed as just one more attempt to extend the reach of mechanical philosophy to a particular area of scientific inquiry.

Having discussed activity liberalism and phenomenal liberalism, let us now direct our attention to the third form of liberalism: social liberalism. The main advantage of social liberalism is that it allows non-human social agents

(e.g., social robots) to qualify as the *social* components of a socio-technical mechanism. For the sake of illustration, imagine a case where a community of social robots coordinate their activities by communicating with one another over an IP-compatible network. Inasmuch as we accept the idea that social robots ought to be regarded as social agents, and we also discover that these agents are components in an integrated mechanism, then we will have met the conditions associated with the mechanistic view, and we can thus conclude that the robots form part of a larger system that ought to be regarded as a social machine.

Whether or not we should be this permissive about the inclusion of non-human agents into a socio-technical mechanism is undoubtedly an issue for further work. (Some may think, for example, that the notion of a social component is best reserved for human individuals.) For present purposes, what matters is that a commitment to social liberalism provides the basis for theoretical integration between the mechanistic view and the sociable machines view. According to the sociable machines view, recall, a social machine is a technological system that possesses the sort of characteristics that secures its status as a “sociable service component” (Meira et al. 2011). By itself, of course, this will not be enough to confirm the status of the technological system as a social machine, since the claim, according to the mechanistic view, is that social agents are the *components* of socio-technical mechanisms; it is not the claim that a social agent is, by itself, a social machine. As things stand, then, the sociable machines view is likely to require some tweaking if it is to be fully compatible with the mechanistic view. In particular, we need to move away from the idea that a social machine is a single social agent (human or otherwise) towards the idea that a social machine is a system whose phenomena are realized by the causal interactions between a collection of social agents (human or otherwise). Inasmuch as proponents of the sociable machines view are prepared to accept this modification, then there seems little reason to regard the mechanistic view as in any sort of tension with the sociable machines view. Social liberalism thus provides the basis for theoretical integration between the mechanistic view and the sociable machines view.

5 Explaining Stories: Causal Narratives and Developmental Explanations

We have seen how a commitment to three kinds of liberalism (activity, phenomenal, and social liberalism) enables the mechanistic view to accommodate a number of existing views of social machines. This, we suggest, provides us with an important opportunity for theoretical integration.

Other views, however, are not so easily accommodated by the mechanistic view. One such view is the narratological view championed by Tarte et al. (2014). Tarte et al.’s characterization of the narratological view is complex, and the exegesis is, therefore, delicate. Nevertheless, the general idea seems to be that by approaching social machines from the standpoint of a storytelling approach,

we are sensitized to the methodological virtues of narratological techniques as part of our attempt to understand social machines. Such methods, Tarte et al. (2014) suggest, are to be commended on the grounds that they help to reveal the social aspects of social machines, aspects that are, according to Tarte et al., at risk of being overlooked or occluded by other approaches.

The first thing to note here is that the mechanistic view is completely neutral as regards the kinds of methods and techniques that are used to discover, analyze, understand, and/or build socio-technical mechanisms. The mechanistic view simply identifies the kind of mechanisms that ought to be targeted for the purpose of identifying social machines; it does not tell us anything about the kinds of methods that ought to be used to discover (and/or implement) such mechanisms. No doubt the selective interest in (online) socio-technical mechanisms does entail some methodological peculiarities. Perhaps, for example, the scale and complexity of socio-technical mechanisms is sufficient to require the use of methods and techniques that are distinct from those used in other scientific disciplines. This would no doubt have important implications for mechanical philosophy, since the majority of debates in mechanical philosophy (and, more broadly, the philosophy of science) have centered on the epistemic practices of scientists working with a limited range of mechanistic examples (see Zednik 2015). Such implications are, to be sure, worth noting, and they should undoubtedly be explored in further work; for present purposes, however, it is important to note that nothing about the mechanistic view imposes constraints on the sorts of methods that ought to be applied to socio-technical mechanisms. In addition to the forms of liberalism discussed above, we thus encounter a form of liberalism that might be referred to as *methodological liberalism*. Courtesy of this particular form of liberalism, the mechanistic view is nicely poised to accommodate the convictions of those who call for methodologically-circumscribed approaches to understanding social machines.

From a methodological standpoint, then, there is no particular tension or incompatibility between the narratological view and the mechanistic view. There is, however, a further point that is worth mentioning here. This concerns the way in which narratological constructs (e.g., stories, narratives, plots, and so on) have been incorporated into theories of mechanistic explanation. Of particular interest, in this respect, is work by Glennan (2010) and Little (2018), both of whom direct their attention to the notion of historical explanation. According to Glennan and Little, historical explanation can be thought of as a particular form of mechanistic explanation; it is a form of mechanistic explanation in which historical events (or, more generally, historical phenomena) are explained courtesy of an appeal to historical mechanisms. Crucially, for our purposes, narratives play an important role as both the targets of historiographic analysis and as the vehicle through which historical explanations are communicated. “Historical explanations,” Little (2018) notes, “typically take the form of causal narratives linking actions, circumstances, and consequences leading up to the outcome to be explained” (p. 414).

The upshot is a better understanding of the relationship between the mechanistic view and the narratological view. Rather than presenting us with a

view that is radically at odds with the mechanistic view, the narratological view emerges as perfectly consistent with the overall effort to chart the structure, organization, and operation of socio-technical mechanisms. In attempting to provide a mechanistic explanation of some online phenomenon, we will undoubtedly need to understand something about the causal interactions between a potentially disparate array of material objects. There is no reason why each of these material objects (human or otherwise) should not be seen as having something akin to its own ‘story’, a story that describes the object’s informational encounters and transactions with other objects. There is, in addition, no reason why these individual (component-level) narratives should not be seen as forming part of a larger narrative structure (just as the narrative structure of a film or novel is, on occasion, formed from an interacting nexus of individual storylines). From the standpoint of the mechanistic view, the ‘story’ of a socio-technical mechanism is likely to be an account that details the way in which a collection of individual components work together so as to realize some phenomenon of interest. As such, it will be a story that emerges from the complex web of interactions between the components of a mechanism. As with the notion of historical explanation, what we encounter here is a form of mechanistic explanation that appeals to the concepts and methods used in narratological analysis. In particular, it is a form of mechanistic explanation that helps us understand something about the structure and dynamics of a socio-technical mechanism. It does this by detailing the way in which a collection of component-level ‘tales’ are spun together so as to form a larger narrative fabric.

The appeal to mechanistic explanation thus provides a means of helping us understand the relationship between the narratological view and the mechanistic view. It also, we suggest, helps us understand the relationship between the mechanistic view and the socially-situated machines view (see Section 2). The socially-situated machines view, recall, regards the human social environment as playing a *causal* role in the development of machine-based cognitive capabilities. In essence, the socially-situated machines view sees a social machine as a purely technological system, albeit one whose (cognitive) capabilities arise as a result of the system’s developmental immersion in a human social environment.

From the standpoint of theoretical unification, the socially-situated machines view presents us with a significant challenge. This is due to the way in which a social machine is conceptualized as a purely technological system under the socially-situated machines view. The mechanistic view, recall, regards a social machine as a system whose phenomena are explained via an appeal to socio-technical mechanisms. There is, however, no reason why a purely technological system (e.g., a machine learning system) is required to rely on such mechanisms. Many human cognitive capabilities (e.g., our capacity for mathematical cognition) are no doubt ones that we acquire as a result of our interactions with a multiplicity of human individuals. But this does not

mean that the *exercise* of such capabilities are *always* ones that depend (in a constitutive sense) on the deployment of socio-technical mechanisms.⁷

The upshot, it seems, is that the mechanistic view is *not* compatible with the socially-situated machines view. Before we accept this conclusion, however, we ought to consider the variety of forms of mechanistic explanation that have been discussed in the philosophical literature. We have already seen the value of this approach in the foregoing discussion of the narratological view. In that case, a consideration of historical explanation (a particular form of mechanistic explanation) turned out to be of crucial importance, helping us establish a conceptual link between two ostensibly distinct approaches to understanding social machines. When it comes to the socially-situated machines view, a similar sort of conceptual link may arise in the context of so-called *developmental explanations* (Parkkinen 2014; Ylikoski 2013). Such explanations are deemed to combine both constitutive and causal mechanistic explanations (see Ylikoski 2013), and they are typically used to explain the developmental emergence of some phenomenon of interest, such as a cognitive capability. This dovetails with the idea that a social machine is, according to the socially-situated machines view, a (technological) system whose cognitive capabilities emerge as a consequence of the system’s developmental immersion in a socially-rich environment. Perhaps, then, the mechanisms that are responsible (in a causal sense) for the emergence of machine-based cognitive capabilities might themselves be taken as evidence for the presence of a social machine. In other words, while we should reject the idea that a socially-situated machine is itself a social machine, the fact that such a system comes into being as the result of a socio-technical mechanism might suggest that, under certain circumstances, a socially-situated machine is produced by a system that does (under the mechanistic view) count as a *bona fide* social machine.

6 Teleological Talk: Goals, Purposes, and Processes

Within the social machine literature, one often encounters talk of “goals,” “purposes,” and “teleology” (see Shadbolt et al. 2019, chap. 1). De Roure et al. (2013), for example, suggest that:

... purpose may be a key notion to help identify individual Social Machines in composed systems. . . Machines and users might simultaneously be members of multiple/different social machines at the same time, and our suggestion is that each social machine is delimited by purpose. (De Roure et al. 2013, pp. 1–2)

One of the virtues of the mechanistic view is that it helps us make sense of this appeal to goal-related (or teleological) talk within the social machine

⁷ We can, at times, perform mathematical calculations in our heads, even if our capacity to perform those calculations is acquired as a result of our prior social interactions with other individuals. In this case, a capacity to perform mathematical calculations might be acquired as the result of a social mechanism, but the actual process of implementing the calculation (the exercise of the capacity) might require nothing more than a neural mechanism.

literature. One of the central tenets of mechanical philosophy is what is sometimes referred to as *Glennan's Law* (Craver 2007a, 2013; Glennan 1996). This is the idea that mechanisms are interest-relative, in the sense that they are individuated according to the phenomena we seek to explain, or, in the case of engineering efforts, the phenomena we seek to implement. In this sense, mechanisms are always *for* something. We say that something is a component of a mechanism courtesy of the fact that it is constitutively relevant to a particular (explanandum) phenomenon. But in the absence of this phenomenon, we cannot determine what the components of a mechanism are. It is this attempt to pin down the explanandum phenomenon that, we suggest, lies at the root of teleological talk within the social machine literature. In essence, to state that a social machine must have a goal or purpose is simply to recognize the crucial role of some explanandum phenomenon in individuating a socio-technical mechanism.

From the standpoint of the mechanistic view, then, there is nothing that is particularly surprising about the persistent presence of teleological talk within the social machine literature. Nor is there any reason to regard teleological talk as in any way inconsistent with the mechanistic view. In fact, such talk emerges as an ineliminable aspect of the attempt to discover (and in some cases, engineer) social machines. Much the same point has been made by Craver (2013), who talks of the ineliminable teleological loading of mechanistic explanations in the life sciences:

...the identification of functions is a crucial step in the discovery of mechanisms. We no longer speak of mechanisms simpliciter, but rather as mechanisms *for* some behavior. Mechanistic descriptions thus come loaded with teleological content concerning the role, goal, purpose, or preferred behavior of the mechanism. This teleological loading cannot be reduced to features of the causal structure of the [world], but it is ineliminable from our physiological, and particularly neural, sciences, precisely because their central goal is to make the busy and buzzing confusion of complex systems intelligible and, in some cases, usable. (Craver 2013, p. 155, original emphasis)

The upshot of all this is a better understanding of teleological talk in the social machine literature. Consistent with the intuitions of many, teleological talk *does* turn out to be important to our capacity to discover (or implement) social machines. Crucially, however, this is not because social machines are required to have a particular purpose, goal, or function, or that they ought to be evaluated relative to the extent to which such purposes, goals, or functions are actually fulfilled. According to the mechanistic view, the status of some system as a social machine is confirmed by the presence of a socio-technical mechanism, and socio-technical mechanisms (as with *all* mechanisms) are individuated relative to some explanandum phenomenon (i.e., some phenomenon that we deem to be of explanatory importance). It is thus the explanandum phenomenon that is crucial to our capacity to discover a socio-technical mechanism and thus identify a social machine. In the absence of this target explanandum

phenomenon—which identifies the goal of a particular explanatory effort rather than the goal of a particular social machine—we will never be in a position to discover a socio-technical mechanism (and thus detect the presence of a social machine).

In addition to helping us make sense of teleological talk, the appeal to Glennan’s Law also helps us accommodate the process view of social machines within the overarching theoretical framework of the mechanistic view. The process view, recall, claims that social machines ought to be regarded as tasks or processes. Such a view might be seen to be in some tension with the mechanistic view, since, according to the mechanistic view, social machines are physical systems rather than tasks and processes. Relative to the foregoing discussion, however, it should now be clear that social machines are not simply physical systems; rather, they are *functionally-individuated* physical systems, i.e., systems that are identified courtesy of an attempt to explain (or, in the case of engineering efforts, to implement) a particular phenomenon. Here, the appeal to processes makes perfect sense, because, from a metaphysical standpoint, processes are one of the kinds of occurrent phenomena that are deemed to be targeted by (constitutive) mechanistic explanations (Kaiser and Krickel 2017). Accordingly, the appeal to processes in the context of the process view can be seen as an attempt to isolate phenomena that (as per Glennan’s Law) are crucial to our capacity to individuate mechanisms. All that the process view adds to this existing notion is the idea that, in the specific case of socio-technical mechanisms, the relevant phenomena ought to be constituted by the joint activity of social and technological components, or, as Tinati and Carr (2012) put it, “any task that requires the co-constitutional involvement of humans and technologies is a form of social machine” (p. 975). This is an idea that is in perfect accord with the mechanistic view. Indeed, inasmuch as we deem the processes targeted by the process view to be those that are constituted or realized by socio-technical mechanisms, then there ought to be no conflict or tension between the proponents of the process view and the mechanistic view. From the standpoint of the mechanistic view, the processes alluded to by the process view are simply processes that are realized by socio-technical mechanisms. It is perfectly possible, then, that the proponents of both the process view and the mechanistic view will direct their attention to the same kinds of processes as part of their attempt to identify social machines.

7 Problem Solving Systems

Thus far, we have seen how the mechanistic view provides opportunities for theoretical integration in respect of many of the views discussed in Section 2. The only view that remains to be discussed is the problem-solving view. According to this view, social machines ought to be regarded as problem-solving systems. Hooper et al. (2016), for example, refer to social machines as “problem-solving entities” and suggest that one of the issues we should address when we encounter a social machine is to ask ourselves “what problem does it solve?”

The problem-solving view presents little in the way of a problem for the mechanistic view. A problem-solving process is clearly a process of sorts,⁸ and we have already seen how a commitment to phenomenal liberalism enables us to accommodate a multiplicity of processes within the mechanistic view (see Section 4). There is, in addition, no need to worry about the apparent need to determine the specific problem that is solved by a social machine, for this probably reflects nothing more than a tacit recognition of the importance of Glennan's Law (see Section 6).

8 Interdisciplinary Connections and Epistemic Practice

The main virtue of the mechanistic view is that it provides us with an opportunity for theoretical integration. In particular, by imposing constraints on different aspects of the mechanistic account, we are able to see how the various views of social machines discussed in Section 2 might be accommodated within an overarching theoretical framework. In respect of the cognitive systems view, for example, the constraints are applied to the phenomena targeted by a mechanistically-oriented explanatory account. If we restrict these phenomena to those of the cognitive variety (e.g., cognitive processes), then we end up with an account that approximates that proposed by proponents of the cognitive systems view. Similarly, if we impose constraints on the activities performed by human individuals and computational systems, then we end up with an account that approximates that proposed by proponents of the content creation view.

The virtues of the mechanistic view do not end with theoretical integration, however. In addition to providing us with a means of understanding social machines from a conceptual standpoint, the mechanistic view also provides us with a particular perspective on the nature of the social machine research effort. While mechanistic explanation is not the only form of explanation employed by those who study the Web and Internet, the appeal to mechanistic explanation does, we suggest, yield a number of important benefits for the science of social machines, as well as the broader discipline of WAIS.

One such benefit concerns the nature of interdisciplinary links with the social scientific community. The science of social machines is often glossed as an interdisciplinary endeavor, one that benefits from input and collaboration with the social sciences. At present, however, it is somewhat difficult to say why such interdisciplinary collaboration is required, or what kind of explanatory effort is

⁸ It should be noted that Hooper et al. (2016) do not provide a definition of problem-solving processes; nor do they suggest how problem-solving processes might be distinguished from processes of the non-problem-solving variety. For present purposes, we are assuming that problem-solving processes—of the sort suggested by the problem-solving view—are simply one of the phenomena (or object-involving occurrents) targeted by constitutive mechanistic explanations. This is consistent with way that Kaiser and Krickel (2017) refer to the phenomena of constitutive mechanistic explanations. Such phenomena, they suggest are object-involving occurrents, where “the occurrents in which objects can be involved are of different kinds; they are, for instance, behaviours, events, *processes*, activities, or states that are manifestations of dispositions” (Kaiser and Krickel 2017, p. 770; emphasis added).

best supported by this collaboration. The mechanistic view provides us with a potential answer to these questions. In particular, it provides the conceptual foundation for interdisciplinary work with mechanistically-oriented approaches in the social sciences—approaches that seek to explain social phenomena by detailing the (social) mechanisms responsible for those phenomena (Hedström 2005; Hedström and Ylikoski 2010). The basis for interdisciplinary collaboration, in this case, stems from a common commitment to mechanistically-oriented modes of explanation and the (partially) social nature of the mechanisms that lie at heart of the mechanistic view (i.e., *socio-technical* mechanisms). Social scientists are, of course, interested in a wide variety of phenomena, and not all these phenomena require an explanatory account that is rooted in the operation of socio-technical mechanisms. Nevertheless, the increasing penetration of the Web and Internet into multiple areas of social life, as well as the emergence of prominent cyber-social phenomena (e.g., cyber-victimization, cyber-crime, etc.), speaks to the likely importance of socio-technical mechanisms in formulating mechanistic explanations of social phenomena. A nice example of this sort of explanatory effort stems from the work of González-Bailón et al. (2014) who explored the role of Twitter in the Spanish Indignados movement—a digitally-coordinated political protest that emerged in May 2011.

One of the benefits of mechanical philosophy is that it helps us understand the epistemic practices of scientists working across a broad array of disciplines. This raises a question about the extent to which the mechanistic view informs our understanding of the epistemic practices employed by those who study the Web and Internet. Answering this question is complicated by a number of factors.

The first complication stems from the fact that WAIS is both an analytic and synthetic discipline. That is to say, research efforts within WAIS are directed to both the analysis of existing systems and the implementation (synthesis) of new systems. By itself, this particular point does not pose much of a problem for the proponent of the mechanistic view. The mechanistic view simply draws attention to the nature of the mechanisms that will need to be supported by an up-and-running system. In this respect, the mechanistic view appears consistent with at least some development-oriented efforts within the social machine community. A nice example of this stems from the work of Murray-Rust et al. (2018) who describe the use of so-called sociograms (diagrams constructed using a specialized modelling language) to construct models of a prospective social machine (see also Shadbolt et al. 2019, p. 70–74). As indicated by the term “sociogram,” these diagrams invariably depict the interactions between a multiplicity of social actors, with some diagrammatic elements (nodes) representing “actors in the system, whether computational or human” and others (edges) representing “the interactions between these actors” (Murray-Rust et al. 2018, p. 767). While Murray-Rust et al. (2018) do not explicitly refer to sociograms as representing socio-technical mechanisms, it should be clear that there is at least a rough correspondence between the diagrammatic elements of sociograms and the graphical depiction of mechanisms

within mechanical philosophy (recall the discussion of Craver diagrams in Section 3).

A second complication relates to the methodological practices of the WAIS community. Unlike many of the scientific disciplines that animate discussions in mechanical philosophy, experimental techniques are something of a rarity in the WAIS literature (even more so when it comes to the literature on social machines). This does not mean that the appeal to mechanistic explanation is inappropriate, however. In studying social machines, researchers frequently rely on the Application Programming Interfaces (APIs) provided by specific social machines. These provide access to voluminous bodies of data about the operation of a social machine. There is, as far as we can tell, no reason why such data cannot be used to formulate mechanistic explanations of a particular phenomenon. Indeed, the data emanating from APIs often provides a useful record of the activities performed by particular entities as part of the run-time operation of a social machine. Inasmuch as this information sheds light on the operation of a socio-technical mechanism, then social machine practitioners may have little need for experimental interventions: experimental interventions are, in effect, rendered unnecessary as a result of the particular bodies of empirical data that are available to the social machine community. Clearly, this is an issue that requires further exploration; nevertheless, there seems little reason to think that the epistemic practices employed by the WAIS community are congenitally ill-equipped to support the process of mechanism discovery. In support of this idea, Zednik (2015) highlights the role played by non-experimental techniques (e.g., computer simulations) in supporting the process of mechanism discovery (see also Kästner 2017). Spring et al. (2017) also to refer to this issue as part of their effort to apply mechanical philosophy to the realm of information and cyber-security.⁹

A final complication arises in respect of the nature of explanatory efforts within WAIS. The problem here is that the discipline, as a whole, shows no commitment to a particular form of explanation. While some explanatory efforts within WAIS are geared to the discovery and description of mechanisms, others refer to the properties of networks and are probably better characterized as attempts at topological explanation (Huneman 2010).¹⁰ As with the more general body of WAIS research, the science of social machines expresses a robust interest in the topological properties of networks. But this does not mean that the science of social machines is dominated by a commitment to topological (as opposed to mechanistic) explanation. Indeed, in many cases, the topological properties of a network are not so much the basis for an explanation of a phenomenon as they are a description of the phenomenon to be explained. In

⁹ Indeed, many of the issues tackled by Spring et al. (2017) (see also Spring and Hatleback 2017) resemble those associated with the effort to apply mechanical philosophy to WAIS. In this respect, we see an important opportunity for collaboration, both with respect to the application of mechanical philosophy to different areas of computer science, as well as the more practical effort to study mechanisms that straddle the social and technological realms.

¹⁰ Such diversity is, of course, not specific to WAIS. As noted by Huneman (2010; 2018), one finds evidence of topological explanation in the biological sciences.

short, the goal of scientific activity in this area is not just to document the types of networks that are to be found in the online realm, it is also to understand something about the *mechanisms* that give rise to these networks (see, for example, Capocci et al. 2006). In this respect, the explanatory interests and concerns of the social machine community intersect with those who embrace a mechanistically-oriented approach to the explanation of social phenomena—the structure of social networks being one of the explananda targeted by such approaches (see Hedström and Ylikoski 2010, p. 58). The only difference is that whereas social scientists direct their attention to mechanisms that are (for the most part) constituted by human individuals and their interactions with one another, the science of social machines is interested in mechanisms that subtend both the social and technological realms. This, at least, is the picture that emerges once we approach the science of social machines from the standpoint of the mechanistic view.

9 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 mechanical philosophy. This account—the mechanistic view—is intended to progress our conceptual understanding of social machines. 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 both social (e.g., human) and technological (e.g., computational) components.

The mechanistic view is intended to move us closer towards an understanding of social machines, providing us, perhaps, with an initial glimpse as to what the shape of an overarching theoretical account of social machines might look like. It does this by drawing on analytic resources that were originally developed to provide a theoretical account of mechanistic explanation within the philosophy of science. In this sense, the present paper highlights the importance of philosophical work relative to the theoretical and empirical agenda of an emerging scientific discipline. Beyond this, however, the present paper focuses attention on a specific class of mechanisms, namely, socio-technical mechanisms. While such mechanisms have seldom been at the forefront of debates in mechanical philosophy, they are, we suggest, an important focus area for future philosophical work.

Acknowledgements This work is supported by the UK EPSRC as part of the PE-TRAS National Centre of Excellence for IoT Systems Cybersecurity under Grant Number EP/S035362/1. We are grateful to two anonymous reviewers for feedback in respect of an earlier draft of the present paper.

References

- Ahlers D, Driscoll P, Löfström E, Krogstie J, Wyckmans A (2016) Understanding Smart Cities as Social Machines. In: 4th International Workshop on the Theory and Practice of Social Machines, Montreal, Canada
- Baumgartner M, Gebharder A (2016) Constitutive relevance, mutual manipulability, and fat-handedness. *The British Journal for the Philosophy of Science* 67(3):731–756
- Baumgartner M, Casini L, Krickel B (2020) Horizontal surgicality and mechanistic constitution. *Erkenntnis* 85:417–340
- Bechtel W (2008) *Mental Mechanisms: Philosophical Perspectives on Cognitive Neuroscience*. Lawrence Erlbaum Associates, New York, New York, USA
- Berners-Lee T, Fischetti M (1999) *Weaving the Web: The Original Design and Ultimate Destiny of the World Wide Web*. Harper Collins, New York, New York, USA
- Berners-Lee T, Hall W, Hendler J, Shadbolt N, Weitzner DJ (2006a) Creating a Science of the Web. *Science* 313(5788):769–771
- Berners-Lee T, Hall W, Hendler JA, O’Hara K, Shadbolt N, Weitzner DJ (2006b) A framework for Web Science. *Foundations and Trends in Web Science* 1(1):1–130
- Burégio V, Meira S, Rosa N (2013a) Social machines: A unified paradigm to describe social Web-oriented systems. In: 1st International Workshop on the Theory and Practice of Social Machines, Rio de Janeiro, Brazil
- Burégio VA, Meira SR, Rosa NS, Garcia VC (2013b) Moving towards “relationship-aware” applications and services: A social machine-oriented approach. In: 6th International Workshop on Evolutionary Business Processes, Vancouver, British Columbia, Canada
- Capocci A, Servedio V, Colaiori F, Burriol L, Donato D, Leonardi S, Caldarelli G (2006) Preferential attachment in the growth of social networks: The internet encyclopedia Wikipedia. *Physical Review E* 74(3):36116
- Clarke B, Leuridan B, Williamson J (2014) Modelling mechanisms with causal cycles. *Synthese* 191(8):1651–1681
- 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*. Clarendon Press, Oxford, UK
- Craver C, Tabery J (2016) Mechanisms in science. In: Zalta EN (ed) *The Stanford Encyclopedia of Philosophy*, Spring 2016 edn, Stanford University, Stanford, California, USA
- Craver CF (2013) Functions and Mechanisms: A Perspectivalist View. In: Huneman P (ed) *Functions: Selection and Mechanisms*, vol 363, Springer, Dordrecht, The Netherlands, pp 133–158
- Craver CF, Darden L (2013) *In Search of Mechanisms: Discoveries Across the Life Sciences*. The University of Chicago Press, Chicago, Illinois, USA
- De Roure D, Hooper C, Meredith-Lobay M, Page K, Tarte S, Cruickshank D, De Roure C (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
- De Roure D, Hendler JA, James D, Nurmikko-Fuller T, Van Kleek M, Willcox P (2019) Towards a Cyberphysical Web Science: A Social Machines Perspective on Pokémon GO! In: Boldi P, Welles BF, Kinder-Kurlanda K, Wilson C, Peters I, Meira WJ (eds) 10th ACM Conference on Web Science, ACM, Boston, Massachusetts, USA, pp 65–69
- van Eck D (2018) Mechanisms and engineering science. In: Glennan S, Illari PM (eds) *The Routledge Handbook of Mechanisms and Mechanical Philosophy*, Routledge, New York, New York, USA, pp 447–461
- Glennan S (1996) Mechanisms and the nature of causation. *Erkenntnis* 44(1):49–71
- Glennan S (2010) Ephemeral mechanisms and historical explanation. *Erkenntnis* 72(2):251–266
- Glennan S (2017) *The New Mechanical Philosophy*. Oxford University Press, Oxford, UK
- Glennan S, Illari PM (eds) (2018) *The Routledge Handbook of Mechanisms and Mechanical Philosophy*. Routledge, New York, New York, USA

- González-Bailón S, Borge-Holthoefery J, Morenoy Y (2014) Online networks and the diffusion of protest. In: Manzo G (ed) *Analytical Sociology: Actions and Networks*, John Wiley & Sons, Ltd, Chichester, West Sussex, UK
- Hall W, Tiropanis T (2012) Web Evolution and Web Science. *Computer Networks* 56:3859–3865
- Harinen T (2018) Mutual manipulability and causal inbetweenness. *Synthese* 195(1):35–54
- Hatleback E, Spring JM (2014) Exploring a mechanistic approach to experimentation in computing. *Philosophy & Technology* 27(3):441–459
- Hedström P (2005) *Dissecting the Social: On the Principles of Analytical Sociology*. Cambridge University Press, Cambridge, UK
- 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, Mulvehill AM (2016) *Social Machines: The Coming Collision of Artificial Intelligence, Social Networking, and Humanity*. Apress, New York, New York, USA
- Hooper C, Bailey B, Glaser H, Hendler J (2016) Social machines in practice: Solutions, stakeholders and scopes. In: Nejd W, Hall W, Parigi P, Staab S (eds) *8th International ACM Web Science Conference*, ACM, Hannover, Germany, pp 156–160
- Horsman DC (2015) Abstraction/Representation Theory for heterotic physical computing. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 373:20140224
- Huneman P (2010) Topological explanations and robustness in biological sciences. *Synthese* 177(2):213–245
- Huneman P (2018) Diversifying the picture of explanations in biological sciences: ways of combining topology with mechanisms. *Synthese* 195(1):115–146
- Illari PM, Williamson J (2012) What is a mechanism? Thinking about mechanisms across the sciences. *European Journal for Philosophy of Science* 2(1):119–135
- Kaiser MI (2018) The components and boundaries of mechanisms. In: Glennan S, Illari PM (eds) *The Routledge Handbook of Mechanisms and Mechanical Philosophy*, Routledge, New York, New York, USA, pp 116–130
- Kaiser MI, Krickel B (2017) The metaphysics of constitutive mechanistic phenomena. *The British Journal for the Philosophy of Science* 68(3):745–779
- Kästner L (2017) *Philosophy of Cognitive Neuroscience: Causal Explanations, Mechanisms, and Experimental Manipulations*. De Gruyter, Berlin, Germany
- Little D (2018) Disaggregating historical explanation: The move to social mechanisms. In: Glennan S, Illari PM (eds) *The Routledge Handbook of Mechanisms and Mechanical Philosophy*, Routledge, New York, New York, USA, pp 413–422
- Luczak-Roesch M, Tinati R, O’Hara K, Shadbolt N (2015) Socio-technical computation. In: *18th ACM Conference on Computer-Supported Cooperative Work and Social Computing*, Vancouver, British Columbia, Canada
- Luczak-Roesch M, Tinati R, Aljaloud S, Hall W, Shadbolt N (2016) A universal socio-technical computing machine. In: *16th International Conference on Web Engineering (ICWE’16)*, Lugano, Switzerland
- Machamer P, Darden L, Craver CF (2000) Thinking about mechanisms. *Philosophy of Science* 67(1):1–25
- Martin U (2016) Computational logic and the social. *Journal of Logic and Computation* 26(2):467–477
- Meira SRL, Burégio VAA, Nascimento LM, Figueiredo E, Neto M, Encarnação B, Garcia VC (2011) The Emerging Web of Social Machines. In: *35th Annual Computer Software and Applications Conference (COMPSAC)*, Munich, Germany
- Menzies P (2012) The causal structure of mechanisms. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences* 43(4):796–805
- Murray-Rust D, Davoust A, Papapanagiotou P, Manataki A, Van Kleek M, Shadbolt N, Robertson D (2018) Towards executable representations of social machines. In: Chapman P, Stapleton G, Moktefi A, Perez-Kriz S, Bellucci F (eds) *International Conference on Theory and Application of Diagrams*, Springer, Edinburgh, UK, Lecture Notes in

- Artificial Intelligence, vol 10871, pp 765–769
- O’Hara K (2012) Trust in social machines: The challenges. In: AISB/IACAP World Congress 2012: Social Computing, Social Cognition, Social Networks and Multiagent Systems, Birmingham, UK
- Palermos SO (2017) Social Machines: A Philosophical Engineering. *Phenomenology and the Cognitive Sciences* 16(5):953–978
- Parkkinen VP (2014) Developmental Explanation. In: Galavotti MC, Dieks D, Gonzalez WJ, Hartmann S, Uebel T, Weber M (eds) *New Directions in the Philosophy of Science*, Springer, London, UK
- Salmon WC (1984) *Scientific Explanation and the Causal Structure of the World*. Princeton University Press, Princeton, New Jersey, USA
- Shadbolt NR (2013) Knowledge acquisition and the rise of social machines. *International Journal of Human–Computer Studies* 71(2):200–205
- Shadbolt NR, Smith DA, Simperl E, Van Kleek M, Yang Y, Hall W (2013) Towards a classification framework for social machines. In: 1st International Workshop on the Theory and Practice of Social Machines, Rio de Janeiro, Brazil
- Shadbolt NR, O’Hara K, De Roure D, Hall W (2019) *The Theory and Practice of Social Machines*. Lecture Notes in Social Networks, Springer, Cham, Switzerland
- Smart PR (2018) Knowledge Machines. *The Knowledge Engineering Review* 33:e11
- Smart PR, Madaan A (2017) The Social Scaffolding of Machine Intelligence. *International Journal on Advances in Intelligent Systems* 10(3&4):261–279
- Smart PR, Madaan A, Hall W (2019) Where the Smart Things Are: Social Machines and the Internet of Things. *Phenomenology and the Cognitive Sciences* 18(3):551–575
- Spring JM, Hatleback E (2017) Thinking about intrusion kill chains as mechanisms. *Journal of Cybersecurity* 3(3):185–197
- Spring JM, Illari P (2019) Building general knowledge of mechanisms in information security. *Philosophy & Technology* 32(4):627–659
- Spring JM, Moore T, Pym D (2017) Practicing a science of security: A philosophy of science perspective. In: *Proceedings of the 2017 New Security Paradigms Workshop*, ACM, Santa Cruz, California, USA, pp 1–18
- Strohmaier M (2013) A few thoughts on engineering social machines. In: 1st International Workshop on the Theory and Practice of Social Machines, Rio de Janeiro, Brazil
- Strohmaier M, Wagner C (2014) Computational social science for the World Wide Web. *IEEE Intelligent Systems* 29(5):84–88
- Tarte S, De Roure D, Willcox P (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
- Tinati R, Carr L (2012) Understanding social machines. In: *ASE/IEEE International Conference on Social Computing and International Conference on Privacy, Security, Risk and Trust*, Amsterdam, The Netherlands
- Tiropanis T, Hall W, Crowcroft J, Contractor N, Tassiulas L (2015) Network Science, Web Science, and Internet Science. *Communications of the ACM* 58(8):76–82
- Wilson RA, Craver CF (2007) Realization: Metaphysical and Scientific Perspectives. In: Thagard P (ed) *Philosophy of Psychology and Cognitive Science*, North-Holland, Oxford, UK, pp 81–104
- Ylikoski P (2013) Causal and constitutive explanation compared. *Erkenntnis* 78(2):277–297
- Ylikoski P (2018) Social mechanisms. In: Glennan S, Illari PM (eds) *The Routledge Handbook of Mechanisms and Mechanical Philosophy*, Routledge, New York, New York, USA, pp 401–412
- Zednik C (2015) Heuristics, Descriptions, and the Scope of Mechanistic Explanation. In: Braillard PA, Malaterre C (eds) *Explanation in Biology: An Enquiry into the Diversity of Explanatory Patterns in the Life Sciences, History, Philosophy and Theory of the Life Sciences*, Springer, Dordrecht, The Netherlands, pp 295–318
- Zednik C (2018) Mechanisms in cognitive science. In: Glennan S, Illari PM (eds) *The Routledge Handbook of Mechanisms and Mechanical Philosophy*, Routledge, New York, New York, USA, pp 389–400