

Foundations and Trends[®] in Web Science

Minds Online: The Interface between Web Science, Cognitive Science and the Philosophy of Mind

Suggested Citation: Paul Smart, Robert Clowes and Richard Heersmink (2017), “Minds Online: The Interface between Web Science, Cognitive Science and the Philosophy of Mind”, Foundations and Trends[®] in Web Science: Vol. 6, No. 1-2, pp 1–232. DOI: 10.1561/1800000026.

Paul Smart

University of Southampton
ps02v@ecs.soton.ac.uk

Robert Clowes

Universidade Nova de Lisboa
robert.clowes@gmail.com

Richard Heersmink

Macquarie University
richard.heersmink@gmail.com

This article may be used only for the purpose of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval.

now

the essence of knowledge

Boston — Delft

Contents

1	Introduction	3
2	Cognition, Cognitive Science and Cognitive Ecologies	7
2.1	What is Cognition?	7
2.2	Cognitive Science, Philosophy and the Web	12
2.3	Varieties of Cognition	16
2.4	The Web as a Cognitive Ecology	30
3	Embodied Cognition	31
3.1	Embodied Interaction	32
3.2	Shaping the Body	33
3.3	Corporeal Extension	36
3.4	The Wearable Web	40
3.5	Measurable Me	42
3.6	The Quantified Cognizer	46
4	Extended Cognition	48
4.1	Web-Extended Minds	49
4.2	Trusting The Web	59
4.3	Dimensions of Cognitive Integration	64
4.4	Extended Memory	68
4.5	Extended Knowers	75

4.6	Supersized Knowers?	79
4.7	Privacy Implications of the Web-Extended Mind	83
5	Embedded Cognition	86
5.1	The Cognitive Commons	87
5.2	The Real World Web	89
5.3	Programmable Worlds	93
5.4	Embed with the Web?	98
5.5	Extended vs. Embedded: Does It Matter?	100
6	Social Cognition	111
6.1	Person Perception	113
6.2	Social Bots	117
6.3	Social Influence	119
6.4	Social Brains, Extended Minds	124
7	Collective Cognition	129
7.1	Socio-Cognitive Circuits	130
7.2	Social Feedback	135
7.3	Collaborative Construction	136
7.4	Social Machines	141
7.5	Knowledge Engines	143
7.6	Leaning on the (Social) Environment	149
7.7	Mandevillian Intelligence	155
8	Extending the Butterfly	159
8.1	Cognitive Connections	161
8.2	Boundary Dissolution	164
8.3	Augmented Cognition?	167
8.4	Web Epistemology	169
8.5	Ecological Engineering	172
9	Conclusion	177
	Acknowledgements	179
	References	180

Minds Online: The Interface between Web Science, Cognitive Science and the Philosophy of Mind

Paul Smart¹, Robert Clowes² and Richard Heersmink³

¹*University of Southampton, ps02v@ecs.soton.ac.uk*

²*Universidade Nova de Lisboa; robert.clowes@gmail.com*

³*Macquarie University; richard.heersmink@gmail.com*

ABSTRACT

Alongside existing research into the social, political and economic impacts of the Web, there is a need to study the Web from a cognitive and epistemic perspective. This is particularly so as new and emerging technologies alter the nature of our interactive engagements with the Web, transforming the extent to which our thoughts and actions are shaped by the online environment. Situated and ecological approaches to cognition are relevant to understanding the cognitive significance of the Web because of the emphasis they place on forces and factors that reside at the level of agent–world interactions. In particular, by adopting a situated or ecological approach to cognition, we are able to assess the significance of the Web from the perspective of research into embodied, extended, embedded, social and collective cognition. The results of this analysis help to reshape the interdisciplinary configuration of Web Science, expanding its theoretical and

empirical remit to include the disciplines of both cognitive science and the philosophy of mind.

1

Introduction

Ever since its inception as a specialized area of scientific enquiry, Web Science has been conceived as a discipline that benefits from (and perhaps depends upon) various forms of interdisciplinary engagement [44, 45, 235, 428]. Although it is perhaps easy to think of the World Wide Web (hereafter simply the Web) as a purely technological system—one which involves the use of protocols, formalisms and software algorithms to support the exchange of information across a global computational network—there is clearly a sense in which the Web is more than just a collection of technological elements. In particular, the Web is as much a social phenomenon as it is a technological one,¹ and this alters the theoretical and empirical remit of Web Science, expanding its scope to include topics that fall within the orbit of the social, psychological and

¹It is perhaps more appropriate to see the Web as a socio-technical phenomenon, rather than a purely social or technological one. From the perspective of Web Science, the Web is typically conceived as a socio-technical system: one that features complex forms of interaction and inter-dependence between individuals, technology and society [368]. Web Science recognizes that technological innovation on the Web may impact society, perhaps giving rise to new forms of social interaction and engagement. However, it also recognizes that technological innovation occurs against a sociological and psychological backdrop: a backdrop that includes, among other things, moral and ethical codes, legal constraints, social conventions and human cognitive capabilities.

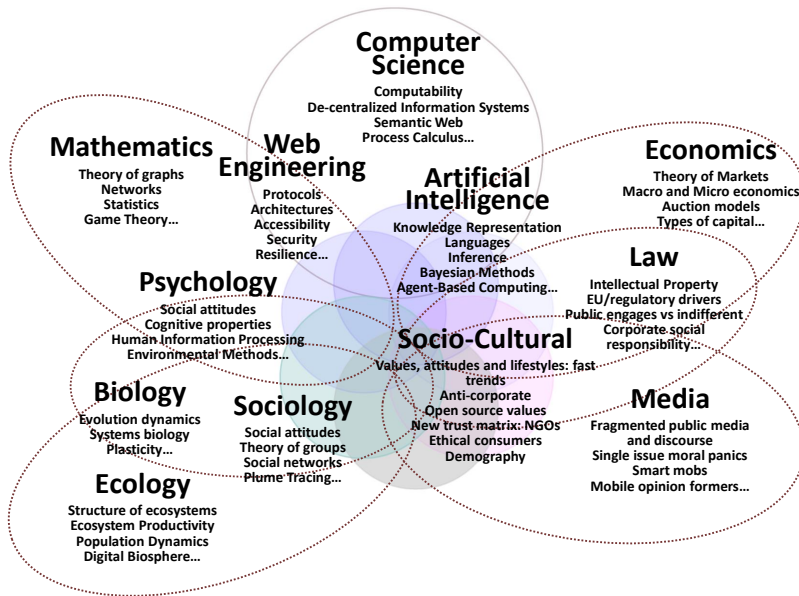


Figure 1.1: The Web Science Butterfly Diagram (<http://www.webscience.org/>).

cultural sciences. Some insight into the interdisciplinary nature of Web Science is provided by the so-called Web Science Butterfly Diagram, which represents an early attempt to map out those areas where the interests and concerns of Web Science converge with those of other disciplines (see Figure 1.1).

A particular set of disciplines are typically seen as relevant to Web Science. Writing in 2008, for example, Hendler et al. [235] suggest that Web Science establishes contact with the disciplines of mathematics, computer science, artificial intelligence, sociology, psychology, biology and economics. Interestingly, philosophy and cognitive science—the disciplinary targets of the present review—are absent from this list. This does not, of course, mean that philosophy and cognitive science are irrelevant to Web Science. Indeed, Hendler et al. [235] encourage researchers to expand the disciplinary reach of Web Science by explicating the links with other disciplines. This is, in fact, one of the aims of the present review. In particular, we seek to highlight areas where

the interests of cognitively-minded scientists and philosophers start to converge with those who seek to study and understand the Web.

From a cognitive science perspective, the Web raises important issues concerning the nature of both human and machine intelligence. A particular focus of interest relates to the effect of the Web on human cognitive capabilities [81], especially those that relate to mnemonic functioning [472], reading [174], and social cognition [160]. Another point of interest concerns the Web's ability to support the emergence of hybrid information processing ensembles that combine the capabilities of both humans and machines. This is where the recent interest in a number of Web-based socio-cognitive phenomena—e.g., collective intelligence [322], collective problem-solving [350], collective sensemaking [284] and collective creativity [557]—start to converge with well-established areas of cognitive scientific research, especially with the study of distributed cognitive systems [254].

In addition to the points of contact with cognitive science, the Web has also begun to attract the attention of the philosophical community [215, 216, 217, 382]. This is particularly apparent when it comes to those areas of philosophy that deal with mental or cognitive phenomena, such as the philosophy of mind and the philosophical study of knowledge (i.e., epistemology) [314, 447, 450, 452]. Work in these areas is relevant to our conceptual understanding of the Web, especially as it relates to our status as cognitive and epistemic agents. At the present time, the epistemic implications of the Web are very much in the media spotlight, with terms such as “post-truth,” “fake news,” “filter bubbles,” and “alternative facts” dominating public discourse [see 390]. The cognitive implications of the Web have also been the subject of much debate, with prominent public figures voicing concern about the effect of the Web on cognitive performance [81, 206]. In view of the potential policy implications of these debates, it is imperative that we have a better understanding of the nature of the relationship between the Web and that wondrous (albeit beguiling) organ that we refer to as the human mind. In particular, it is important that the scope of current debates be broadened to accommodate the wealth of work that has been undertaken to improve our conceptual understanding of cognitive phenomena.

The present review represents an attempt to heed the call of the early Web Science pioneers by expanding the interdisciplinary scope of Web Science. In essence, we seek to expand the reach of the butterfly's wings (see Figure 1.1) to accommodate the disciplines of cognitive science and the philosophy of mind. As is evidenced by the present review, there is a substantial literature to support this expansionist agenda. Given the centrality of cognition to our species-specific capabilities, as well as the level of public and scientific interest in the Web, now is arguably a good time to review this literature and explicate the nature of the linkages that connect the science of the Web to the sciences of the mind.

2

Cognition, Cognitive Science and Cognitive Ecologies

2.1 What is Cognition?

The main goal of cognitive science is to advance our understanding of cognition and the mind. For the most part, the same is true of the philosophy of mind. Despite some rather obvious methodological differences between the disciplines, the goal of cognitive science and philosophy of mind remains the same: to improve our understanding of cognition and the mind. But what is it, exactly, that lies within the intellectual cross hairs of these disciplines? What exactly is cognition? What is the mind? And why should Web Science be concerned with either of them?

These are important questions, the discussion of which would no doubt provide the basis for a monograph (or monographs) several times the length of the one you are now reading. For the moment, we will focus our attention on the first (and undoubtedly much more vexed) issue of what is meant by the terms “mind” and “cognition.”¹ (The relevance of cognition to Web Science is something that will become

¹For the most part, the present review talks of “cognition” more than it does the “mind.” We treat the mind as that entity or system that performs or engages in cognition.

clearer during the course of the present review, although something of a cursory response can be found in Section 2.2.)

So, what is cognition? Unfortunately, there is, as yet, no definitive answer to this question. This does not mean that philosophers and cognitive scientists (as well as the rest of us) have no understanding of what cognition is. Neither does it mean that philosophers and cognitive scientists have no idea as to what cognitive scientists should be studying. It simply means that there is, as yet, no commonly agreed linguistic formula for specifying what it is that makes something a *bona fide* cognitive process, or (more broadly) what it is that makes something a *bona fide* cognitive phenomenon.

Within philosophical circles, the problem of identifying the characteristic features of cognitive phenomena (e.g., the things that make a particular process a cognitive process) is referred to as the attempt to resolve the “mark of the cognitive” [1, 2, 3, 4, 5]. The aim, in essence, is to determine what it is that enables us to subsume things like memory, attention, problem-solving, perception, and so on, under the cognitive umbrella. As is noted by Adams and Garrison [3], a mark of the cognitive would enable us to determine what it is that makes a seemingly diverse set of processes instances of the same type—that is cognitive:

It is easy to give a list of cognitive processes. They are things like learning, memory, concept formation, reasoning, maybe emotion, and so on. It is not easy to say, of these things that are called cognitive, *what makes them so*. Knowing the answer is one very important reason to be interested in the mark of the cognitive. [3, p. 340, original emphasis]

There have been many important and interesting attempts to discover the mark of the cognitive. Adams and Aizawa, for example, attempt to limn the realm of the cognitive by appealing to the idea that “cognition is constituted by certain sorts of causal processes that involve nonderived content” [2, p. 67]. Such a view leads Adams and Aizawa to countenance a form of contingent intracranialism about the (human?) mind. They suggest that only neural processes, as a matter of contingent fact, are able to give rise to cognitive phenomena.

Other attempts to resolve the mark of the cognitive include accounts that appeal to reason-guided behaviors [3], characteristic sets of causal processes [1], and the presence of conscious experiences [see 152]. Unfortunately, none of these proposals has received widespread philosophical assent. The problem, in most cases, is either that a specific proposal places too many restrictions on cognition (e.g., it equates cognition with consciousness), or that it appeals to concepts that are no less enigmatic than the notion of cognition itself [see 8]. Consider, for example, the following attempt by Mark Rowlands [414, pp. 110–111] to specify a set of criteria that are deemed to be jointly sufficient in determining whether a particular process P should be counted as a cognitive process:

1. P involves *information processing*—the manipulation and transformation of information-bearing structures.
2. This information processing has the *proper function of making available* either to the subject or to subsequent processing operations information that was, prior to this processing, unavailable.
3. This information is made available by way of the production, in the subject of P , of a *representational state*.
4. P is a process that *belongs* to the *subject* of that *representational state*.

As is noted by Adams and Garrison [3], none of these criteria are entirely unproblematic. What, for example, is the precise meaning of the term “representational state?” What counts as a “proper function?” And what does it mean to say that a process “belongs to” the subject that is the bearer of a representational state? None of these questions, it seems, are likely to yield a straightforward answer (or at least none are likely to yield an answer that would be unequivocally accepted by the philosophical community).

In response to the lack of philosophical progress concerning the mark of the cognitive, some theorists have questioned the need to define the term “cognition.” Allen [8], for example, advocates “a pluralistic stance towards uses of the term ‘cognition’ that eschews the urge to

treat cognition as a metaphysically well-defined ‘natural’ kind” (p. 1). Cognitive science, Allen suggests, requires a concept of cognition that is able to accommodate the diversity of systems that are the focus of cognitive scientific investigations. In view of this diversity, Allen suggests that there is little reason to assume that the realm of the cognitive can be neatly delineated with respect to essentialist criteria. Neither is there any reason, he argues, for cognitive scientists to be overly worried about the absence of a discipline-fixing definition that determines the scope of their scientific efforts:

Cognitive scientists should no more be embarrassed about their lack of an all-purpose definition or categorical description of cognition than biologists are about their inability to define “life,” “species” or any other number of terms, or for physicists to be able to give a unified account of such fundamental notions as “force” and “matter.” [8, p. 14]

For the purposes of the present review, we will embrace a view of cognition that emphasizes its role in the genesis and organization of intelligent behavior. This is, in fact, a view that dates back to the earliest days of cognitive science [see 223]. It is also a view that has managed to survive the upheaval caused by seismological shifts in the scientific and philosophical landscape over the past several decades. As noted by Hatfield [223], the notion that “cognition is information processing that explains intelligent behaviour” (p. 361) is one that has been largely unaffected by traditional dialectics, such as the distinction between symbolic vs. sub-symbolic processing and the representational vs. non-representational character of cognition. Such a view is also able to accommodate the sort of distinctions that have been made regarding what might be called ‘cognitive kinds’, i.e., particular forms of cognition (see Section 2.3). (This is of particular relevance in the context of the present review; for much of the review is based around a particular subset of these cognitive kinds, e.g., extended, embodied and collective cognition). In support of this compatibility, Andy Clark, one of the leading proponents of a particular form of cognition known as extended cognition, suggests that perhaps the best way to conceive of

cognition is to take a “cognition is as cognition does” approach. In this case, issues of cognitive status are fixed by an attempt to account for certain kinds of behavior. “The notion of a cognitive process,” Clark suggests, “is best unpacked as the notion of a process that supports certain kinds of behavior (actual and counterfactual)” [104, p. 93]. A similar position is adopted by another affiliate of the extended cognition camp: Michael Wheeler [542]. He too regards cognition as something that is individuated with respect to certain kinds of behavior:

...behaving appropriately (e.g., adaptively, in a context-sensitive fashion) with respect to some (usually) external target state of affairs, should be counted as displays of intelligence and as outcomes of cognitive processing. [542, p. 3]

This, then, provides us with a view of cognition that we will refer to as the *intelligent behavior view* of cognition:

Cognition (Intelligent Behavior View)

Cognition refers to whatever it is that gives rise to intelligent behavior. What makes (e.g.) a process a cognitive process is the role that the process plays with respect to the generation of intelligent behavior, where the notion of *intelligent behavior* is to be understood as behavior that is appropriate, adaptive, flexible and coordinated with respect to environmental and organismic circumstances (i.e., context-sensitive).

This is the view of cognition that we will adopt throughout the remainder of the review. It is a view that is sufficiently generic to accommodate the various kinds of cognition that we will encounter in subsequent sections. It is also, it should be clear, a view that is largely agnostic with respect to the metaphysical character of the ‘things’ that generate (or give rise to) intelligent behavior. This is important, because it allows for differences in the explanatory accounts that are delivered by cognitive science (e.g., accounts that explain behavior in terms of cognitive *processes* or cognitive *mechanisms*). Some of these differences are no doubt merely idiomatic, and thus of no substantive scientific relevance. Others, however, seem to be

more important. As will become clear in a subsequent section (see Section 2.3), cognitive science is a discipline that, at a practical level, focuses its attention on the physical mechanisms that are responsible for the expression of intelligent behavior. The focus of interest, in this case, concerns the material constituents of mechanisms (e.g., neurons) and their patterns of interaction (e.g., the propagation of bioelectric signals). Such mechanisms are deemed to be *cognitive mechanisms* whenever they support the expression of intelligent behavior. But not everyone speaks in terms of physical mechanisms. Sometimes, behavior is explained with respect to the presence of cognitive states and processes, and it is far from clear that the explanatory referents of such accounts (i.e., states and processes) are the same as those targeted by mechanism-based accounts. Such ostensible differences in explanatory focus are easily accommodated by the intelligent behavior view of cognition. What makes a process (or state) a *cognitive* process (or *cognitive* state) is thus the role that such processes (or states) play with respect to the generation of intelligent behavior. Similarly, when it comes to mechanism-based explanations, what makes a mechanism a *cognitive* mechanism is simply the role that such mechanisms play with respect to the generation of intelligent behavior.

2.2 Cognitive Science, Philosophy and the Web

Relative to the intelligent behavior view of cognition (see Section 2.1), we can see cognitive science as, in a general sense, the science that deals with the forces and factors that give rise to intelligent behavior. This is broadly consistent with the way cognitive science has been conceived by philosophers and cognitive scientists. Bechtel et al. [35], for example, offer the following definition of cognitive science:

Cognitive science is the multidisciplinary scientific study of cognition and its role in intelligent agency. It examines what cognition is, what it does, and how it works. [35, p. 3]

As with Web Science, cognitive science is a multidisciplinary endeavor, involving contact with disciplines such as philosophy, psychology,

Artificial Intelligence (AI), neuroscience, linguistics, and anthropology. Such disciplines serve as the basis for what might be thought of as the cognitive science counterpart to the Web Sciences Butterfly Diagram: the *Cognitive Science Hexagram* (see Figure 2.1).




Figure 2.1: The Cognitive Science Hexagram [see 352].

Another point of commonality between cognitive science and Web Science, aside from their multidisciplinary outlook, concerns their status as analytic and synthetic disciplines. Both Web Science and cognitive science are clearly disciplines that seek to analyze existing systems of one sort or another, either Web-based or cognitive (and sometimes both!). In addition to this, however, Web Science and cognitive science are both disciplines that are, at least in part, concerned with the synthesis (i.e., development) of *de novo* systems (again Web-based, cognitive and sometimes both). This is particularly apparent in the case of Web Science, which as a specialized area of computer science, is suitably poised to extend the Web's existing technological and computational fabric. As is noted by Berners-Lee et al. [45]:

Physical science is an analytic discipline that aims to find laws that generate or explain observed phenomena; computer science is predominantly (though not exclusively) synthetic, in that formalisms and algorithms are created in order to support particular desired behaviour. Web science has to be a merging of these two paradigms; the Web needs to be studied and understood, and it needs to be engineered. [45, p. 3]

For the most part, cognitive science is an analytic discipline: it studies a rich array of naturally-occurring (biologically-based) cognitive systems. Beyond this, however, cognitive science is also concerned with the engineering of cognitive systems, either for the purposes of improving our understanding of existing systems,² or for the purposes of building *de novo* forms of cognitive organization that are able to emulate (and sometimes surpass) the capabilities of biologically-based systems (such as ourselves). It is in respect of this synthetic orientation that we encounter one of the important overlaps between Web Science and cognitive science. For the current preoccupation with machine intelligence and intelligent systems engineering is one that is, to an ever-greater extent, of mutual interest to both disciplines. Some of the more notable points of convergence, in this case, come with respect to the use of Semantic Web technologies to support advanced forms of machine reasoning [46], the use of the Web or parts thereof (e.g., Wikipedia) as the informational substrate for cognitive computing platforms [171], and the capacity of the Web (especially the Social Web) to support the emergence of novel forms of machine intelligence [451]. There is, in addition, a rather obvious point of overlap when it comes to the development of Web-based systems that are intended to yield collective intelligence [322, 350], or which aim to augment the cognitive capabilities of individual Web users [337]. All these areas touch on issues of philosophical relevance. Should we, for example, see the representational commitments of the Semantic Web as violating a “claim of intrinsic suitability” [see 102, p. 107] regarding the sorts of material substrate that are sufficient for intelligent behavior? Do the derivative semantics of online symbolic

²A classic example of this sort of approach is provided by Braitenberg in [64].

encodings rule out the possibility of genuine machine cognition [1]? Does the virtual nature of the online environment exclude the possibility of cognitively-relevant forms of material embodiment [544]? Does the global and social nature of the Web support the social scaffolding of machine intelligence [463]? And does our philosophical conception of human knowledge influence our approach to the design of certain kinds of Web-based system [382]?

When it comes to their status as synthetic disciplines, therefore, there are a number of points of contact between Web Science, cognitive science and the philosophy of mind. In addition, to this, however, a degree of convergence can also be found in respect of a number of more analytically-oriented issues and concerns. This is particularly apparent when it comes to the cognitive implications of the Web. Does, for example, the nature of our interaction with the Web yield a general enhancement in cognitive functioning, or is it the case that exposure to the Web leads to a form of cognitive diminishment? Is the Web a cognitive boon (to be embraced), or is it a cognitive burden (to be endured)? It turns out that these questions are much more problematic than we might have thought. And the problem is not merely one of running suitably controlled experiments. Instead, the problem lies in the way we conceptualize cognition; in particular, the way we think about the human mind and its relationship to the physical, social and technological environments in which we humans are materially situated. Here, then, is a further point of contact with the philosophy of mind: in order to get an appropriate understanding of the cognitive significance of the Web—the implications that the Web has for human cognizing—we will need to understand the various ways in which the human mind is situated in the “causal structure of the world” [see 422]. As will become clear in the next section, the results of this philosophical analysis provide us with rich array of conceptual positions. When it comes to questions about cognition and the Web, the answer we get will probably depend on the kind of conceptual lens we choose to don.

2.3 Varieties of Cognition

Even so much as a cursory glance at the cognitive scientific literature will reveal that cognition comes in a rich variety of flavors. In addition to the forms of cognition associated with specific biological taxa (e.g., avian cognition, primate cognition, human cognition, and so on), the cognitive science literature is littered with terms and concepts referring to ‘cognitive kinds’ of a somewhat more esoteric nature. These include extended cognition [7, 102, 108], embedded cognition [417, 418], scaffolded cognition [16, 483], embodied cognition [11, 432, 433, 434, 435], situated cognition [409], distributed cognition [240, 254, 255], group cognition [381, 498, 499], social cognition [288], enactive cognition [141], grounded cognition [28, 29, 386], augmented cognition [477], metacognition [396], and so on. Cognition, it seems, is a many varied (or at least a multi-faceted) thing, and the medley of cognitive flavors on offer seems to be of sufficient richness as to rival the offerings of even the most cosmopolitan of Italian gelaterias.³ How should we make sense of this bewildering array of cognition-related concepts? And what, moreover, is the significance of this profusion of cognitive kinds for the discipline of Web Science?

The present review is, in part, an attempt to answer these questions. While we do not seek to cover every kind of cognition on offer, the distinction between various forms of cognition is, we suggest, of both conceptual and methodological significance. For the purposes of this review, we limit our attention to the following forms of cognition: embodied cognition (see Section 3), extended cognition (see Section 4), embedded cognition (see Section 5), social cognition (see Section 6), and collective cognition (see Section 7).

³Just to complicate the picture, cognitive scientists typically make a distinction between different forms of cognitive processing, such as attention, memory, perception, learning, problem-solving, decision-making, and so on [see 165]. Typically, each of these cognitive process *types* is subject to further forms of conceptual decomposition. In the case of memory, for example, cognitive psychological research is typically organized with respect to different kinds of memory. These include, *inter alia*, episodic, prospective, working, autobiographical, emotional, semantic, implicit, explicit, long-term, short-term, and procedural memory [see 20].

As a means of helping us understand these particular forms of cognition, it will help to view cognition from a number of different viewpoints, perspectives or dimensions. Ultimately, such dimensions may serve as the foundation for a conceptual framework that is able to accommodate many of the cognitive kinds that appear in the contemporary cognitive scientific literature. For the time being, however, we will limit our attention to just those dimensions that are relevant to the kinds of cognition covered by the present review. These dimensions are as follows:

- **Agent Nature:** The nature of the agent that participates in cognitive processing (e.g., human vs. machine cognition).
- **Agent Cardinality:** The number of agents involved in cognitive processing (e.g., individual vs. collective cognition).
- **Cognitive Domain:** The domain over which cognitive processing routines operate (e.g., social vs. mathematical cognition).
- **Cognitive Mechanisms:** The properties (including spatial location and material composition) of the mechanisms that realize cognitive processing (e.g., extended vs. embedded cognition).

As a means of introducing the specific forms of cognition that are covered in the remainder of the review, it will be useful to examine these dimensions in a little more detail.

2.3.1 Agent Nature

The nature of the agent that participates in cognitive processing is perhaps the most straightforward way to understand cognitive kinds. As mentioned above, biological criteria sometimes serve as the basis for specialized forms of cognitive scientific effort, and this gives rise to forms of cognition centered on particular species (or collections thereof). Much of contemporary cognitive science is, of course, concerned with human cognition (i.e., the kind of cognition exhibited by human agents). But other kinds of cognition are also of significant interest. Prominent examples, in this respect, include avian cognition [86], primate cognition [506], cephalopod cognition [137], and even plant cognition [73].

In the present review, we will be concerned exclusively with human cognition. We therefore seek to understand the points of contact between Web Science and the realm of human cognitive processes. This does not mean that other kinds of (agent-nature-oriented) cognition are irrelevant to Web Science. To appreciate this, we need only reflect on the growing importance of machine intelligence and machine learning to the Web Science community [451, 452]. This particular kind of cognition (i.e., machine cognition) is easily discerned once we direct our attention to the *nature* of the agent that lies at the heart of some episode of cognitive processing.

2.3.2 Agent Cardinality

In the attempt to understand the forces and factors responsible for the genesis and organization of intelligent behavior, it is mostly the behavior of individual agents (e.g., individual human subjects) that is the primary focus of explanatory interest and attention. In addition to these *individual* forms of cognition, however, cognitive science has also been preoccupied with cognitive processes involving *multiple* individuals. Research in this area goes by a variety of names, such as group cognition [497, 498, 499, 500], team cognition [128, 129], swarm cognition [509], distributed cognition [240, 254, 510, 511], and collective cognition [131, 460]. While the term “distributed cognition” is probably the most familiar of these terms, we suggest that the term “collective cognition” is the most appropriate way of referring to cognitive processes involving multiple individuals. Collective cognition, on this view, is an umbrella term that subsumes the notions of distributed, team, group and swarm cognition (see Figure 2.2). It refers to situations in which cognitive processing routines are distributed across multiple independent (cognitive?) agents, irrespective of whether or not additional (non-agential) resources (e.g., technological artefacts) are also involved in the realization of the relevant cognitive processing routines.⁴ Relative

⁴In essence, collective cognition occurs when multiple individuals are involved in cognitive processing routines. The result is that multiple individuals (and sometimes physical resources) are seen to be constitutive elements of a larger cognitive organization or collective cognitive system.

to this conceptualization, we can see group, team, swarm and distributed cognition as specific forms of collective cognition (see Figure 2.2). For the purposes of this review, these sub-types of collective cognition are defined as follows:⁵

- **Distributed Cognition:** A form of collective cognition involving (non-agential) artefactual resources in addition in agents. For the most part, distributed cognitive systems can be thought of as socio-technical systems that engage in collective cognitive processes, e.g., collective problem-solving.
- **Team Cognition:** A form of collective cognition featuring relatively high levels of entitativity.⁶
- **Group Cognition:** A form of collective cognition featuring relatively low levels of entitativity.
- **Swarm Cognition:** A form of collective cognition in which the relevant agents are typically eusocial insects, e.g., ants, bees, termites, etc.

Discussions of collective cognition tend to go hand in hand with discussions of collective intelligence. In fact, each of the forms of collective cognition discussed above can be coupled with a specific form of collective intelligence (Figure 2.2). In general, the term “collective intelligence” is used to refer to the capacity of some collective (i.e., multi-agent) organization to engage in intelligent processing or produce intelligent outcomes [see 322]. In discussing the notion of distributed intelligence, for example, Heylighen [238] suggests that distributed intelligence is the

⁵Three dimensions are relevant to the distinction between different kinds of collective cognition. These include the entitativity of the agent collective (team vs. group cognition), the extent to which inter-agent communication involves artefactual mediation (distributed cognition), and the cognitive sophistication of the agents that are involved in collective cognitive processing (swarm cognition).

⁶Entitativity refers to “the extent to which an assemblage of individuals is perceived to be a group rather than an aggregation of independent, unrelated individuals” [181, p. 15]. In the present context, the term is used to refer to the cohesion of a group, i.e., the strength of the ties that link one member of an agent collective to other members of the same collective.

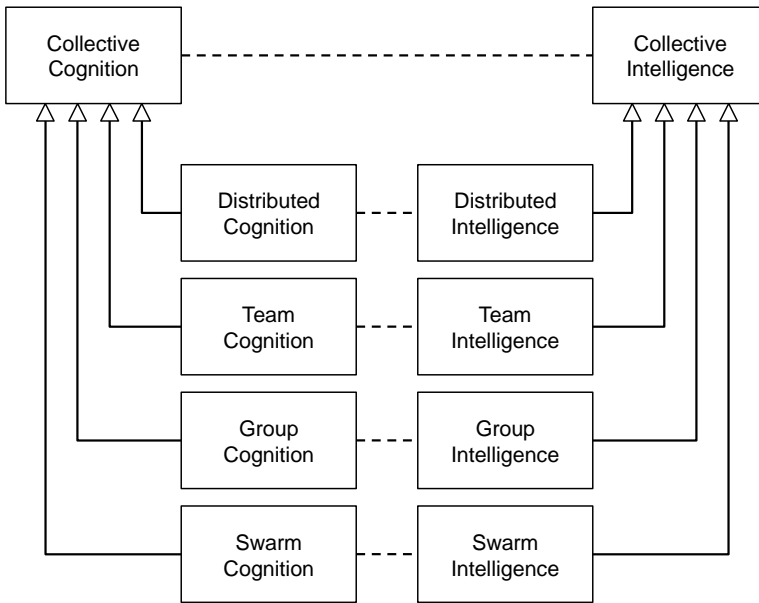


Figure 2.2: A number of forms of collective cognition have been discussed within the philosophical and cognitive scientific literature. These include distributed [254], team [129], swarm [509] and group cognition [381, 474, 498]. A number of forms of collective intelligence have also been discussed. These include distributed [238], team [168], swarm [60, 272], and group intelligence [326]. Each of these forms of collective intelligence is, we suggest, associated with a particular form of collective cognition.

“ability to solve problems collaboratively, by integrating the contributions from a broad assembly of human and technological agents” (p. 500).

In the context of the present review, we are concerned with both individual and collective forms of cognition. Much of the review is devoted to cognition of the individual variety; however, in Section 7 we turn our attention to collective cognition and examine why issues of collective cognition and collective intelligence are of mutual interest to members of the Web Science, cognitive science and philosophy of mind communities.

2.3.3 Cognitive Domain

In addition to looking at cognition from the standpoint of agent nature and agent cardinality, we can also focus our attention on the *domain* of cognitive processing—the kinds of things that cognition is about. From this perspective, we are able to identify a number of forms of cognition, including spatial cognition [147], metacognition [396], mathematical cognition [74], temporal cognition [559], and social cognition [288]. What is common to these forms of cognition is that they refer to cognitive processing that is undertaken with respect to some domain of interest. Spatial cognition, for example, refers to cognitive processes that deal with the processing of spatial information and the generation of behavior that is coordinated with respect to the spatial environment.

From the perspective of Web Science, social cognition is a form of cognition that is of particular interest and relevance. As its name suggests, social cognition refers to cognition that occurs in respect of the spatial domain, e.g., the processing of information derived from the social environment. In particular, social cognitive processes are ones that operate over socially-relevant information, contributing to our ability to make sense of social situations and the behavior of social actors [218, 288]. In the context of the present review, the relevance of such processes to Web Science is discussed in Section 6.⁷

2.3.4 Cognitive Mechanisms

The final dimension we wish to discuss relates to the nature of the mechanisms that are responsible for cognitive processes (or, more generally, cognitive phenomena). There are two points of interest that emerge in relation to such mechanisms. The first concerns the nature of the material elements that make up (or comprise) a cognitive mechanism. The second involves a philosophical distinction between what is

⁷Social cognition is, unfortunately, something that is apt to be confused with the notion of collective cognition. Chi [91], for example, talks of the potential of the Social Web to enhance the ability of groups to remember, think and reason, and he glosses such enhancements as a form of “augmented social cognition.” In fact, according to the approach adopted in the present review, what Chi is really talking about here is a form of collective cognition, i.e., a form of cognition that involves the participatory involvement of multiple human individuals.

called *constitutive mechanistic explanation* and *causal (or etiological) mechanistic explanation*. We will discuss each of these points in turn.

Turning first to the nature of the material elements that make up a mechanism, it should be relatively clear that a significant chunk of cognitive scientific work is concerned with the role of the biological brain in realizing cognitive phenomena. Indeed, the attempt to understand the neural bases of cognitive processes can be seen as one of the central features of disciplines such as cognitive neuroscience, behavioral neuroscience, psychopharmacology and neuropsychology. For these disciplines, the mechanisms of interest are, in most cases, neural mechanisms—they are mechanisms that consist of the cellular elements of the biological brain, i.e., neurons. It is these neural mechanisms that serve as the basis for our (mechanistically-informed) understanding of cognitive phenomena. For the sake of convenience, let us refer to this particular form of cognition as *neurocognition*. Neurocognition is thus the kind of cognition we encounter when the mechanisms that realize a cognitive state or process are comprised of elements drawn solely from the neurological realm.

Now, it might be thought that all forms of biological cognition are also forms of neurocognition; for what else, other than neurons, could be relevant to our understanding of the cognitive processes exhibited by (at least) vertebrate species? In fact, one of the challenges to neurocognition comes from research into what is called embodied cognition [433, 435]. Although there are a number of different views as to what is meant by the term “embodied cognition,” a common feature of embodied cognition research is the emphasis that is placed on extra-neural bodily factors in shaping the course of cognitive processing [11, 433, 434]. Research into embodied cognition thus emphasizes the way in which an organism’s bodily structure or physical actions help to constrain (and sometimes constitute) cognition. A somewhat trivial example is provided by the way in which the placement of an organism’s sensory apparatus (the position of their eyes and ears) helps to structure the incoming sensory array in ways that support perceptual processing [534]. Another example concerns the way in which dynamically evolving

motor state variables can help to guide the expression of adaptive behavioral responses [e.g., 336].

Of course, relative to our intuitions about the role of the biological brain in realizing cognitive phenomena, the notion of embodied cognition can appear unusual (and perhaps implausible). It should be noted, however, that in its most general sense, the term “embodied cognition” is simply a way of highlighting the *relevance* of extra-neural bodily factors to cognitive phenomena. As is discussed below, there are at least two ways of understanding this appeal to the relevance of bodily factors: one contentious; the other, much less so. Let us refer to the more contentious variety as *strong embodied cognition* and the less contentious variety as *weak embodied cognition*.

Focusing first on the weaker form of embodied cognition, embodied cognition theorists sometimes seek to highlight the causal significance of bodily factors with respect to cognitive phenomena. In these cases, the realization base for cognitive phenomena is still located in the neurological realm. In other words, weak embodied cognition views the body as a means to shape or influence the neural mechanisms that realize cognitive phenomena. This is perhaps the best way of approaching a particular strand of embodied cognition research that emphasizes the role of the body in shaping the way we think about the world. Lakoff and Johnson [292], for example, suggest that bodily factors influence the structure of some of our most primitive concepts. Linguistic expressions such as “I am on top of the situation,” “he is under my control,” “she is the head of an organization,” and “I am facing the future,” all seem to rely on spatially-oriented concepts that are tied to the details of our physical embodiment. The body, in this case, is clearly relevant to our understanding of the structure and organization of the human conceptual economy; however, there is nothing particularly radical about this claim. In fact, the claim that the body is exerting a causal influence on neurally-realized cognitive processes need be no more radical or contentious than the claim that a strong cup of coffee influences cognition by affecting aspects of brain biochemistry.

The strong version of embodied cognition differs from its weaker counterpart in insisting that extra-neural bodily elements are not just

exerting a causal influence on cognitive mechanisms; instead, the claim is that such elements are also constituent parts of such mechanisms. The claim, in other words, is that the mechanisms that realize cognitive phenomena include, as literal constituents, the material elements of the extra-neural body. The cognitive mechanisms, in this case, are not solely constituted by neural cells; instead, they extend beyond the neural realm to encompass a range of other physiological resources. If all this sounds highly implausible, it may help to consider the role of non-neural glial cells in modifying aspects of brain function [15, 201, 362, 373, 387]. One view of such cells is that they exert a causal influence on cognitively-relevant neural mechanisms. But given the nature of the interaction between glial and neural cells, is there any principled reason to think that they could not (under any circumstances!) form part of the physical fabric that realizes certain kinds of cognitive phenomena? This seems unlikely, especially since the two cell types are relatively similar, at least with respect to material criteria. But note that once we accept the basic possibility of glial cells forming part of the mechanism that realizes cognitively-relevant phenomena, then the neurocognitive spell is broken; for we are now acknowledging that cognitive phenomena can be realized by mechanisms that are *not* solely constituted by elements drawn from the neural realm. And if we are prepared to accept the possibility of hybrid cognitive mechanisms consisting of both neural cells and glial cells, then why not go further and accept the possibility that other kinds of non-neural corporeal resource might (under certain circumstances) also form part of the mechanistic substrate that realizes cognitive phenomena? This is the point at which issues of embodied cognition start to align themselves with the theoretical and empirical interests of Web Science. For the notion of a non-neural corporeal resource is not something that applies solely to the realm of biological body parts (e.g., your left hand); it also applies to the realm of technological artefacts, such as Web-enabled devices. Inasmuch as we allow for this possibility, then it seems that Web-enabled devices (e.g., a smartphone) could form part of the *bodily* platform that mediates our embodied, sensorimotor engagements with the online world (see Section 3.3).

Now let us turn to another point of interest when it comes to cognitive mechanisms. This point relates to the notion of mechanistic explanation and the distinction between constitutive and causal varieties of mechanistic explanation. The first thing to note here is that when we attempt to understand the mechanistic bases of cognitive phenomena, we are adopting an explanatory approach that goes by the name of mechanistic explanation. Essentially, what are trying to do is understand how a cognitive process is realized by a set of material elements that work together as part of an integrated mechanism. A useful characterization of mechanistic explanation is provided by Bechtel and Richardson [36]. They suggest that:

By calling explanations mechanistic, we are highlighting the fact that they treat the systems as producing a certain behaviour in a manner analogous to that of machines. . . A machine is a composite of interrelated parts. . . that are combined in such a way that each contributes to producing a behaviour of the system. A mechanistic explanation identifies these parts and their organization, showing how the behaviour of the machine is a consequence of the parts and their organization. [36, p. 17]

This notion of mechanistic explanation should be perfectly acceptable to members of the scientific community; for mechanistic explanations are a common feature of scientific practice [133, 134]. Whatever the nature of the phenomenon that is the focus of scientific interest, scientists are typically concerned with the discovery, delineation and description of mechanisms. Indeed, in many cases, mechanistic explanations lie at the heart of our scientific understanding of phenomena. We do not, for example, proclaim to have a complete understanding of biological memory until we have a grasp of the mechanisms that are responsible for memory-related phenomena (e.g., the cellular processes that mediate changes in synaptic strength).

Importantly, mechanistic explanations come in at least two varieties: constitutive mechanistic explanations and causal (or etiological) mechanistic explanations [266]. We have already touched on the distinc-

tion between these two explanatory schemes in the earlier discussion of embodied cognition. Recall that the distinction between the strong and weak versions of embodied cognition turned on the kind of role that non-neural bodily elements were seen to play with respect to the neural mechanisms that realized cognitive phenomena. On the one hand, bodily factors were seen to exert a causal influence on (neurally-realized) mechanisms (i.e., weak embodied cognition); on the other hand, bodily elements were seen to form part of the cognitive mechanisms themselves (i.e., strong embodied cognition). The critical difference between these two forms of embodied cognition actually relates to the distinction between causal and constitutive relevance [132]. When an element is seen to exert a causal influence on the operation of a mechanism, it is deemed to be of causal relevance to the mechanism of interest, and thus of causal relevance to the phenomenon that is the target of mechanistically-oriented explanatory accounts. The result is that when our explanatory focus is oriented towards factors of causal relevance, the kind of mechanistic explanations that we develop are of the causal or etiological variety (i.e., causal mechanistic explanations). Such explanations seek to explicate the causal factors that influence the operation of a mechanism, and they thus help to provide a better understanding of the factors that causally influence phenomena. This kind of mechanistic explanation is distinct from explanations that seek to describe the structure, organization and operation of mechanisms. In this case, we are concerned with constitutive mechanistic explanations, i.e., explanations that tell us how phenomena are linked to the elements that lie internal to the mechanism. This is where issues of constitutive relevance come to the fore: elements that are constitutively relevant to some phenomenon of interest are ones that are components of the mechanism that realizes the phenomenon of interest. These are the elements that are the focus of interest for constitutive mechanistic explanations [see 132].

The distinction between constitutive and causal mechanistic explanations (and constitutive and causal relevance) is useful when it comes to understanding the nature of the philosophical fault lines that separate the proponents of two prominent cognitive kinds: embedded cognition [417, 418] and extended cognition [102, 108]. According to

the proponents of extended cognition, cognition can sometimes extend beyond the biological borders of the human individual to encompass (or incorporate) resources drawn from the extra-organismic environment (i.e., the environment that lies external to the traditional metabolic boundaries of skin and skull). This is essentially a claim about the constitutive relevance of resources that might not otherwise have been deemed to form part of the realization base for cognitive phenomena. At this point, it should be clear that the notion of extended cognition bears a close resemblance to the strong version of embodied cognition discussed above. Strong embodied cognition claims that elements of the extra-neural corporeal environment are constitutively relevant to cognition, and it thus invites us to consider the explanatory relevance of the body from the standpoint of constitutive mechanistic explanations. Something similar can be said about extended cognition. The only notable difference is that extended cognition expands the scope of our explanatory efforts to include elements that lie external to the biological borders of the body. In both cases, we encounter a commitment to what might be called *extended cognitive mechanisms* [see 163, 265, 268]. Such mechanisms, if they exist, provide us with an extended view of cognition: cognitive states and processes (or, more precisely, the mechanisms that realize such states and processes) are seen to be materially-extended, entangling all manner of extra-neural and extra-organismic resources into complex time-variant nexuses of cognitively-relevant information processing. The result, according to some theorists, is that the machinery of the mind is not confined (or at least not always confined) to the intra-cranial realm:

. . . the actual local operations that realize certain forms of human cognizing include inextricable tangles of feedback, feed-forward, and feed-around loops: loops that promiscuously criss-cross the boundaries of brain, body, and world. The local mechanisms of mind, if this is correct, are not all in the head. Cognition leaks out into body and world. [102, p. xxviii]

A significant challenge to extended cognition comes in the form of embedded cognition [417, 418]. The proponents of embedded cognition

challenge the claim that the boundaries of cognition extend beyond the traditional biological borders of the human agent. Extra-organismic resources, they suggest, can influence cognitive processing in sometimes deep and unexpected ways; however, they reject the claim that such resources should be seen as intrinsic to the mechanisms that realize cognitive states and processes. Although cognitive mechanisms can be seen to depend on aspects of the external environment, extra-organismic resources should, according to embedded cognition theorists, be seen to be causally rather than constitutively linked to cognitive processes. From the standpoint of embedded cognition, therefore, the machinery of the human mind remains confined to the intra-cranial realm.

By now it should be clear that the difference between embedded and extended cognition is linked to issues of causal and constitutive relevance. In the case of embedded cognition, extra-organismic elements are seen to exert a causal influence on cognitive processing, and they are thus of causal relevance to cognitive phenomena. This, however, does not make such elements part of the physical fabric (i.e., part of the cognitive mechanism) that realizes cognitive phenomena. Extended cognition, in contrast, views extra-organismic elements in a different way. In this case, the extra-organismic elements are seen to be constitutively relevant to cognitive phenomena. The result is that extended cognition views extra-organismic resources as, on occasion, proper parts of the mechanisms that realize cognitive phenomena.

This way of thinking about embedded and extended cognition is also relevant to issues of embodied cognition. In particular, the distinction between causal and constitutive relevance helps us view embodied cognition from either an embedded or an extended perspective. From an embedded perspective, for example, we can see a bodily element as exerting a causal influence on neural processing. The bodily element, in this case, is deemed to be causally relevant to cognitive processing, and the result is an embedded form of embodied cognition, corresponding to the weak version of embodied cognition discussed above (see Figure 2.3). Alternatively, if we view things from an extended perspective, we can see the bodily elements as forming an integral part of the mechanisms that realize cognitive phenomena. The bodily element, in this case,

	Corporeal Element	Extra-Organismic Element
Causal Relevance	Embedded Embodied Cognition	Embedded Cognition
Constitutive Relevance	Extended Embodied Cognition	Extended Cognition

Figure 2.3: A taxonomy of embedded, extended and embodied cognition, organized with respect to two factors: the object that is the focus of explanatory attention (corporeal vs. extra-organismic elements) and the explanatory relevance of the focal object to cognitive phenomena (causal vs. constitutive relevance). The first factor concerns the focus of explanatory attention—whether our attention is focused on corporeal (i.e., bodily) elements or elements that lie external to the biological borders of the human individual (i.e., extra-organismic elements). The second factor concerns the extent to which the focus of our explanatory interest is causally or constitutively relevant to the phenomena that we seek to explain. If our attention is focused on (extra-neural) corporeal elements, and we think of the causal or constitutive relevance of these elements to cognitive processing, then we get the embedded and extended versions of embodied cognition. If, however, our attention is focused on extra-organismic elements (i.e., elements that lie external to the biological borders of the human individual), and we think of the causal or constitutive relevance of these elements to cognitive processing, then we get the conventional (non-embodied) forms of embedded cognition and extended cognition.

is constitutively relevant to cognitive processing, and the result is an extended form of embodied cognition, corresponding to the strong version of embodied cognition discussed above (see Figure 2.3).

In the context of the present review, the Web is discussed from the perspective of embodied, extended and embedded cognition. Issues of embodied cognition are discussed in Section 3. For the most part, the focus of Section 3 is on the extended (or strong) version of embodied cognition. Section 4 discusses the Web from the standpoint of extended cognition. Here, we explore the extent to which the informational and technological elements of the Web can be seen as part of the material fabric that realizes cognitive phenomena. In other words, Section 4 explores the idea that the Web may, on occasion, be incorporated into extended cognitive mechanisms. Finally, the implications of embedded cognition for our understanding of human–Web interactions are explored in Section 5. This section explores the ways in which the Web can be

seen to exert a causal influence on cognitive phenomena. In this case, there is no commitment to the idea that the Web is an intrinsic part of the mechanisms that realize cognitive phenomena.

2.4 The Web as a Cognitive Ecology

Over the past several decades, the Web has emerged as an important part of the material environment in which human cognition occurs. To an ever-greater extent our cognitive and epistemic endeavors are shaped by the nature of our interactions with the online (i.e., Web-based) world. This view ties in nicely with what is known as an ecological approach to cognition [see 27, 30, 256, 321, 361, 511]. One of the features of this approach is the emphasis that is placed on cognitive ecosystems (i.e., complex networks of material forces and factors that span brain, body and world) as part of the attempt to explain and understand human cognitive capabilities. Hutchins [256], for example, suggests that our attempt to understand “cognitive phenomena must include a consideration of the environments in which cognitive processes develop and operate” (p. 706). From an ecological perspective, therefore, the Web can be seen as part of the ecosystem for human cognition: it serves as an increasingly important part of the material environment in which an ever-expanding array of human cognitive and epistemic activities are situated [see 460].

This sort of ecologically-oriented view of the Web is one that will be adopted throughout the remainder of the review. Not only is such a view the perfect conceptual partner to the various forms of cognition we discuss in subsequent sections (i.e., embodied, extended, embedded, social and collective cognition), it is also a view that helps to focus attention on the nature of our interactions and engagements with the online environment. In particular, an ecological approach enables us to link Web-based forms of social participation with the notion of niche construction [102, 293, 482]—an idea that potentially transforms our understanding of the socio-technical origins and perhaps the ultimate destiny of the human mind (see Section 8.5).

3

Embodied Cognition

At first sight, it might be thought that embodied cognition is of little relevance to Web Science. Work in embodied cognition tends to focus on situations in which we are actively engaged with the real world, exploiting all manner of sensorimotor cues in order to realize intelligent thought and action. The nature of our interaction with the Web seems far removed from this sort of situation. Although we might be justified in seeing the Web as an important part of the context in which cognition occurs—part of the material backdrop against which our thoughts and actions take shape—it is by no means clear that the details of our physical embodiment really matter that much when it comes to understanding the cognitive significance of the Web. As Canny and Paulos [75] note, “cyberspace has been built on Cartesian ideals of a metaphysical separation between mind and body: When we enter cyberspace, even a 3D world, it is the ‘mind’ that enters. . . The body stays outside” (p. 276).

This notion of the environmentally-decoupled and physically-disembodied nature of our online interaction contrasts with the main thrust of theoretical and empirical work in embodied cognitive science. In this case, active, real-time engagement with the external world is a central

element of the embodied cognitive science agenda. The result is that the Web looks to be an unlikely focus of philosophical or scientific interest for those adopting an embodied approach to cognition.

As we shall see, there are a number of reasons to doubt the integrity of this rather pessimistic conclusion. The present section thus attempts to highlight a number of points of contact between Web Science and the science of embodied cognitive systems.

3.1 Embodied Interaction

One reason to think that the Web is relevant to embodied cognition comes from the changing nature of our everyday interactions with the Web. As is noted by Smart [449], the advent of mobile and portable computing technology is progressively altering our sense of what it means to engage with the online world. In place of conventional forms of interaction, in which we interact with the Web via a browser interface while seated at a desktop computer, it is increasingly common for us to engage with the Web as part of our embodied interactions with the wider physical environment. Mobile devices, such as smartphones, thus enable us to interleave our interactions with the Web and the ‘real world’ in a manner that seems to blur the traditional distinction between ‘offline’ and ‘online’ modes of interaction [see 175, 177]. In addition, as new kinds of Web-enabled device become available, so the palette of physical actions and gestures that we use to interact with the Web is expanding. Touchscreens have clearly played an important role, here, with swiping and zooming emerging as more-or-less standard parts of our gestural lexicon. Other kinds of interactivity aim to capitalize on the way in which we typically interact with a common array of physical artefacts and objects, helping to support various forms of *embodied interaction* [see 150] with the online realm. These sorts of innovations help to situate the Web at the heart of our everyday interactions with the world, and they help to make the Web a potent source of interest for those who approach cognition from an embodied perspective.

As a means of reinforcing this particular point, consider the attempt of Matsumoto et al. [332] to develop a Web-enabled umbrella. The umbrella features a variety of sensors—e.g., Global Positioning

System (GPS), compass, accelerometer, etc.—and it is able to project Web-based content directly into the user’s field of view by virtue of a projection device focused on the underside of the umbrella canopy. By providing the user with a range of interactive opportunities (e.g., the normal turning, dipping, and twisting actions that people perform with umbrellas), and by also integrating information from a variety of sensors and Web services, the umbrella is able to present a variety of forms of context-relevant information that take into account both the user’s physical location as well as their current interests and activities. Interestingly, Matsumoto et al. [332] describe their work as part of an effort to realize what they call the *Embodied Web*: a form of enhanced interactivity in which natural embodied interactions are used to interact with the Web and “make our experience in the real world more engaging and active” [332, p. 49].

The introduction of mobile and portable computing devices thus marks an important shift in the way we access the Web. The traditional vision of online interaction is one in which we are sat in front of a desktop computer, accessing the Web through a conventional browser-based interface (such as Internet Explorer or Google Chrome). In these cases, we are encouraged to see the flow of our thoughts as somewhat decoupled from the ‘real world’, as occurring in response to remotely located information resources, and as being largely unaffected by events in the sensory periphery of the computer screen. This vision is not necessarily incorrect; however, it is rapidly being superseded by a different vision—one that situates the Web at the heart of an ever-expanding array of everyday activities.

3.2 Shaping the Body

The technologies we use to access the Web influence the structure of our physical actions. In order to access particular bodies of information from the Web, we thus need to resort to a sequence of actions whose topographic structure is, at least in part, determined by the properties of the device that mediates our contact with the online world. There is nothing particularly profound about this observation; for it should be clear that the nature of our physical interactions with a desktop com-

puter are not the same as those associated with the use of a smartphone. When we think of the various ways in which we access the Web, we can see the properties of our devices as enforcing (or at least promoting) the adoption of different kinds of bodily structure. The resulting vision is perhaps best captured by the metaphor of a lock and key. This sees the process of accessing online information as a form of ‘unlocking’ activity—one in which we adjust our bodies (the key) so as to fit the conformational constraints imposed by a physical device (the lock) in order to gain access to the online world:

Lock-and-Key Metaphor

In retrieving information from the Web, the properties of a device influence the shape of our physical actions. The device can be thought of as a lock and the user’s body as a key. Dynamic adjustments in the shape of the user’s body are required to ‘fit’ the device in order to gain access to online information.

The value of this metaphor is that it highlights the role of device-related properties in altering our physical (bodily) actions as part of our interactive engagements with the Web. Of course, such changes in body-related parameters are precisely the sorts of things that embodied cognitive scientists tend to be interested in. We see this in a range of studies that focus on the relationship between bodily variables and cognitive outcomes. When it comes to body posture, for example, research has revealed that postural factors can influence the recall of autobiographical memories [146] and alter subjective responses to task performance [479]. It is also known that body posture can influence psychosocial processes. In one study, for example, Huang et al. [247] demonstrated that posture could influence the cognitive and behavioral manifestations of power, with an expansive body posture increasing one’s sense of power and activating associated action tendencies.

Facial expressions have also been the focus of considerable research attention by the embodied cognitive science community. Particular interest has been expressed in the facial feedback hypothesis [69]—the idea that facial expressions play a role in regulating overt behavior, as

well as aspects of subjective experience. Requiring subjects to furrow their brow while performing a cognitive task, for example, increases the perceived difficulty of the task [283, 479], and such manipulations can alter the metacognitive processes associated with specific learning experiences [283]. In accord with the facial feedback hypothesis, facial expressions have also been implicated in various emotional responses. In particular, interventions that alter the activation of specific facial muscles have been shown to accentuate certain kinds of emotional response [484]. The facial feedback hypothesis has also been invoked to account for the anti-depressant efficacy of Botulinum toxin (BOTOX[®]) treatment [172, 320, 552]. Given that BOTOX[®] administration reduces the activation of certain facial muscles, the claim is that depressive symptoms are alleviated as the result of a sustained inability to ‘express’ negative emotions (e.g., sadness).

Some of the most important findings regarding the relationship between cognitive and corporeal variables have emerged in relation to gesture. In particular, studies suggest that gesture may play an important role in alleviating cognitive load and supporting learning [125, 126, 192, 193]. Children who gesture while learning a mathematical concept are more likely to maintain what they have learned, as compared to children who do not gesture [125]. The same appears to be true of adults attempting to learn a foreign language [9]. Cook et al. [126] also report that gesture plays a role in mnemonic processing, with gesture at the encoding stage of a memory task boosting subsequent retrieval performance. Crucially, these results are not limited to the case of spontaneous gesture; they also occur with interventions that are deliberately intended to influence gestural movements. Instructed gesture thus yields the same sorts of cognitive benefits as those found in the case of spontaneous gesture, and these benefits apply to a range of cognitive processes, such as working memory [127, 388], mnemonic recall [126] and learning [125].

At this point, it should be clear that the role of technology in shaping the structure of our bodily responses is a topic of crucial importance to embodied cognitive science. As we move from one device (e.g., a desktop computer) to another (e.g., a smartphone), so our

bodies must ‘contort’ to fit the modes of interaction afforded (and sometimes mandated) by the device in question. In the light of the empirical findings concerning posture, facial expression, and (perhaps most importantly) gesture, we can now see that such conformational adjustments in corporeal structure are likely to come with cognitive consequences. This raises a host of issues for the Web Science community, especially for those whose empirical interests are directed to the realm of Human–Computer Interaction (HCI). One focus of research attention relates to the cognitive consequences of using specific kinds of device to access the Web. Does the transition from keyboard/mouse to multi-touch screen, for example, affect cognitive performance as a result of subtle shifts in (e.g.) mnemonic processing and/or the alleviation of cognitive load? And what is the effect of enhanced ‘mobility’ on our creative and communicative capabilities [see 529]?

Another line of enquiry relates to the design of future human–Web interfaces. Can we structure our modes of interaction with the Web so as to exploit body-mediated changes in cognitive performance? Something along these lines is, in fact, proposed by Klemmer et al. [279] who see issues of embodiment as helping to provide new directions for interaction design. Some initial work by Segal [426] is also of interest here. Segal reports on research that aims to explore the extent to which specific ‘gestural interfaces’ can be used to support cognitive processing. He suggests that interfaces supporting the expression of certain gestures can foster the use of representational and computational strategies that are conducive to problem-solving success.¹

3.3 Corporeal Extension

Another way of highlighting the theoretical and empirical significance of embodied cognition to Web Science comes in the form of what might be called *corporeal extension*. This is the idea that Web-enabled

¹The role of technology in supporting productive shifts in psycho-affective well-being is another line of research that is germane to the design of gestural interfaces. Such work is inspired by research into the facial feedback hypothesis. The idea of being forced to smile before one can use a technological device is one expression of this sort of idea [see 357].

devices may, on occasion, function as literal prosthetic extensions of an individual's biological body. Such a claim is, of course, likely to be immediately rejected by those who insist that the notion of a 'body' is grounded in appeals to issues of material composition. It might be claimed, for example, that the body is a *biological* entity and only *biological* systems can be said to qualify as embodied systems.² The legitimacy of this view is called in question, however, as soon as we consider the way in which a number of non-biological elements can serve as replacements for conventional body parts. Consider, for example, the way in which our intuitions about the centrality of bio-material considerations are liable to shift once we reflect on the 'corporeal' status of a variety of non-biological objects, such as tooth implants, prosthetic limbs and cochlear implants. What seems to be important, in these cases, is the way in which some object or set of objects mediates our sensorimotor engagements with the world. Our ears therefore count as part of our body because they assist with the transduction of certain kinds of energetic fluctuation in the ambient environment; our legs count as part of our body because of the way they service our locomotory objectives; and our teeth count as part of our body because of the way they enable us to physically prepare certain kinds of matter for the processes of digestion and absorption.

Another reason to reject the idea that embodied cognitive science should restrict its interests to the realm of biological bodies comes from work in AI [see 93]. It thus seems that some of the core lessons and ideas from embodied cognitive science are just as applicable to the design of synthetic intelligent systems as they are to our understanding of naturally-occurring (biological) systems. Consider, for example, Clark's [97] description of the can-collecting robot, Herbert. Clark describes how certain aspects of the robot's bodily motion (in this case, the rotation of the robot's 'torso') works to achieve an appropriate physical alignment between the robot and a target object—one that helps to secure the success of subsequent reaching movements.

²See Ziemke [562] for an useful overview of the notion of embodiment, as well as a discussion of the kinds of bodies are able to support embodied cognition.

What, then, is the right way to think about the body? One answer to this question comes from approaches that emphasize the functional contribution of body parts in mediating our sensorimotor engagements with the wider extra-corporeal environment [99, 101, 102]. Clark [101] thus suggests that we should identify the body with whatever it is that just so happens to serve as “the locus of willed action, the point of sensorimotor confluence, the gateway to intelligent offloading, and the stable (though not permanently fixed) platform whose features and relations can be relied upon in the computation of certain information-processing solutions” (pp. 55–56). The claim, in essence, is that we should identify the body with whatever it is that fulfils the functional role typically played by the biological body in giving rise to intelligent behavior. This leads to a functionally-oriented conception of the body—one that is fully abstracted from issues of material implementation.³

Inasmuch as we accept a functionally-oriented approach to the body, then it seems that our current arsenal of portable and mobile Web-enabled devices are viable candidates for corporeal extension. In other words, there seems to be no reason (at least in principle) why Web-enabled devices should not count as *candidate* objects that play the role of literal body parts. The crucial question, of course, is to what extent such devices actually are apt for bodily incorporation. It has to be said that very few empirical studies are available that deal directly with this issue; what work is available tends to rely mostly on circumstantial or anecdotal evidence. The mobile phone has probably been the focus of most attention when it comes to issues of corporeal extension [407, 421]. Drain and Strong [151], for example, suggest that the smartphone “. . . becomes incorporated within the assemblage of bodily appendages, environmental features, and artifacts that we encounter in everyday life, to the point where the phone can be considered as a prosthetic

³A functional approach to understanding the body is defended by Wheeler [544]. The advantage of such an approach, Wheeler suggests, is that it allows us to apply the notion of embodied cognition to virtual agents, e.g., synthetic agents that exist solely in the online environment. A similar idea surfaces in respect of recent work into *computational embodied cognition* [468]. In this case, issues of embodied cognition are explored using virtual 3D environments built on top of state-of-the-art computer game engines. Such work builds on earlier work involving the use of virtual environments for the purposes of artificial life research [440, 495].

extension of ourselves” (p. 190). A number of studies have also revealed that users often regard their mobile phones as extensions of their ‘self’ or body [183, 371, 372]. Perhaps such results should not come as a surprise given the way many individuals now relate to their mobile phones. To an ever-greater extent the smartphone is an indispensable instrument that enables the individual to negotiate the various social, cognitive and epistemic challenges that they confront as part of their daily life [151, 353]. This is often reflected in the deep emotional attachments that people have with their mobile devices [353]. As Vincent et al. [524] note “for some people [their mobile device] has become almost an extension of their body as they hold and fondle the device even when the device is not in use” (p. 72).

There are a number of ways in which issues of corporeal extension (especially as they relate to Web-enabled technologies) could be explored in future scientific and philosophical work. These include, but are not necessarily limited to, the following:

- **Functional Properties:** One option is to focus on the nature of our interaction with Web-enabled devices and assess the extent to which such forms of interaction satisfy the sort of functional criteria alluded to by (e.g.) Clark [102]. These criteria emphasize the functional role of a bio-external resource in mediating our sensorimotor interactions with the world.
- **Neurophysiological Changes:** A second option is to explore changes in brain activity using (e.g.) neuroimaging or electroencephalographic (EEG) recording techniques [421]. Traditionally, this has been the preferred approach for investigating tool-related extensions to what is commonly referred to as the *body schema* [259]. The body schema, in this case, is “a suite of neural settings that implicitly (and non-consciously) define a body in terms of its capabilities for action, for example, by defining the extent of ‘near space’ for action programs” [99, p. 272].
- **Action Kinematics:** Changes in the kinematics of action execution are another means of studying the mutability of the body schema [77, 78]. Cardinali et al. [77], for example, show that the

use of a tool induces changes in the kinematics of grasping and pointing movements, in a manner that suggests changes in the somatosensory representation of bodily properties. Interestingly, the results of Cardinali et al. [77] point towards an enduring change in the body schema: one that affects the dynamics of sensorimotor processing even in the absence of the tool. Such results point to something of an ‘after-effect’ or ‘capability echo’ concerning the impact that previously integrated tools have on our action-related capabilities.

- **Perceptual Changes:** Finally, a number of studies have relied on experimental techniques involving subjective changes in so-called after-image phenomena [67, 79, 399]. This research is based on the idea that the assimilation of a tool into the body schema can be indexed by action-consistent changes in tool-related after-images. A key assumption is that the after-image of a grasped tool should fade from view if the tool is released from the hand and falls to the floor.

The importance of future work in this area should not be underestimated. If we accept that Web-enabled devices can function as literal body parts—as prosthetic technological extensions that enable us to sense, manipulate, exploit, and alter the online world—then they can be seen as putting us in direct contact with the online environment (just as our biological body mediates our access to the real, physical environment). Technologies thus provide us with a means to extend the boundaries of the embodied agent to the very ‘edge’ of the online world. By virtue of the corporeal assimilation of Web-enabled technologies and devices, we extend the reach of our perceptuo-motor capabilities, and in doing so, we potentially transform the shape of the human cognitive system [see 276].

3.4 The Wearable Web

Irrespective of whether or not we buy into the notion of corporeal extension, there should be little doubt that the trend of technology development is to yield devices that are able to establish ever-more intimate

forms of contact with our sensorimotor surfaces. As is usefully pointed out by Lynch [318], “we should not just concern ourselves with whether our knowledge is extending ‘out’ to our devices; *our devices are extending in*, and with them, possibly the information they bring” (p. 299, original emphasis). Wearable technologies are, of course, the most obvious example of these ‘inner-reaching’, intimate devices, with smartwatches and activity trackers now widely available. Other forms of wearable technology promise (or perhaps threaten) to encroach even further on the biological body. Notable areas of research in this space includes work on so-called artificial skin (or e-skin) devices [470, 556], wireless contact lenses [307], smart fabrics [180], and neural implants [41, 42, 119].

The use of wearable technologies to track our bodily states and actions is one area where wearable technologies may be of crucial cognitive significance. As wearable technologies become more sophisticated, so they become more capable of sensing our physiological and behavioral states, and this enables them to contextualize their modes of operation to support cognitive processing. Many contemporary devices already feature this kind of context sensitivity. For example, if we are located in New York, and we wish to visit a particular museum, we can rely on the GPS capabilities of our Web-enabled devices—an iPhone let’s say—to contextualize information retrieval operations, and thus deliver information that is likely to be of particular relevance to our task-specific needs and concerns. One example of this sort of context sensitivity is described by Hattori et al. [224]. They propose an activity-based approach to information retrieval, in which information about a user’s location and likely activities are used to refine search queries and contextualize information retrieval operations.

In all likelihood, future wearable technologies will enable us to exploit a far richer range of body-related cues and affordances than is possible with today’s technology. This is important, because by monitoring body-related information it becomes possible to deliver new forms of context sensitivity. Consider, for example, work by Koriat and Nussinson [283] to investigate the physiological correlates of the “feeling of knowing” [see 282]. They report that the tension of the corrugator muscle can be used to detect the subjective experience we have when we feel we know

something. Inasmuch as future devices are able to track physiological signals that index a variety of epistemic feelings (such as the feeling of knowing or the feeling of difficulty) [see 347], it is possible that future technologies may be able to adaptively modify their modes of operation to support human end-users with regard to a variety of cognitive- and epistemically-relevant performances. In fact, something along these lines is proposed by Kunze et al. [289]. They suggest that the use of mobile sensing technologies portends an era in which technology will be able to recognize and monitor various forms of cognitive activity, revolutionizing our understanding of the factors that contribute to optimal cognitive performance, as well as providing new ways for technology to shape and support our cognitive routines.

3.5 Measurable Me

The use of technology to record or track information about individual human subjects is a central element of work that goes under the heading of the *quantified-self* [315, 492].⁴ This is a term that is used to refer to any form of self-tracking activity, where the information that is tracked is typically of a biological or behavioral nature. Current forms of self-tracking thus include the recording of body weight, energy levels, time usage, heart rate, body temperature, exercise patterns, sleep quality, sexual activity, diet, dreams and blood chemistry (see Figure 3.1).

One of the implications of the quantified-self movement is that it provides a greater degree of awareness regarding one's bodily states and processes. Self-tracking technologies are thus sometimes seen as a means of creating a digital dashboard for the biological body, enabling individuals to tap into a wealth of previously inaccessible data. Some writers thus talk about self-tracking devices as supporting the emergence of technological 'exosenses' that extend the reach of the body's sensory capabilities:

... the quantified self may become additionally transformed into the extended exoself [i.e., a suite of exosenses] as data

⁴A variety of other terms are sometimes used to refer to the same phenomenon. These include "lifelogging," "measured me," "self-tracking" and "self-surveillance."



Figure 3.1: A visualization of exercise and sleep data from the Fitbit App (source: <https://www.fitbit.com/uk/app>).

quantification and self-tracking enable the development of new sense capabilities that are not possible with ordinary senses. The individual body becomes a more knowable, calculable, and administrable object through QS [quantified-self] activity, and individuals have an increasingly intimate relationship with data as it mediates the experience of reality. [492, p. 85]

The general idea, therefore, is that self-tracking affords a new way by which body-related information can come to influence the course of cognitive processing. In particular, by virtue of their ability to make body-related information explicitly accessible and perceptible through other senses, the vision of the quantified-self opens the door to forms of embodied cognition in which issues of technological mediation are of central significance. It has long been known, for example, that individual cognitive and emotional responses can be shaped by feedback loops involving physiological signals [e.g., 518]. In view of this, it seems entirely possible that future technologies could play an important role in determining the way in which our thoughts and actions (as well as the thoughts and actions of others) are influenced by body-related information [see 262].

Conceived of as a purely epistemological exercise—i.e., as a means of getting to know the ‘self’—self-tracking appears to be a commendable strategy. Indeed, the use of *reliable* technologies, coupled with the use of cloud-based data analytic techniques and informative visualizations, is an approach that is largely consistent with a number of views in mainstream epistemology, most notably virtue reliabilism [203] and process reliabilism [196]. There is, however, a concern with the way in which self-surveillance provides the basis for a technological transformation of the self. We thus encounter the idea that self-tracking gives rise to another kind of self—a digital self: a self that is digitally created from the various streams of data that emanate from the biological body. A variety of other terms have been used to refer to this notion of a digital self, including data doubles [316, 415], algorithmic identities [89] and digital doppelgängers [58]. The idea, however, is roughly the same in all cases: the use of tracking technologies creates a digital representation that mirrors or reflects the properties of the original self (i.e., the self of the self-tracker).

It might be thought that the digital self is merely a digital rendition of information that is available via the standard, biologically-based routes to self-related knowledge. It seems, however, that things are not quite that simple. The digital self, it is sometimes claimed, reveals ‘hidden’ dimensions of the original self that might not otherwise be available. Crucially, such dimensions are presented in a form that is concrete, tangible and untainted by the vagaries of wishful thinking, fallible memory, and cognitive bias. The result is a somewhat different ‘picture’ of the self. By virtue of numerical quantification and graphical visualization, tracking technologies perhaps reveal something about the ‘real you’—they provide a window into the true self. In this sense, the digital self perhaps becomes more ‘real’ than the original self: our sense of who and what we are is transformed as a result of our attempts to establish a computationally-mediated quantitative grip on the otherwise unruly realms of cognitive, behavioral and physiological flux [see 58].

The result of all this is a complex vision of the relationship between technology and the self. In particular, the self is seen to be something that emerges from the complex forms of reciprocal influence that exist

between the biological body and the surrounding technological environment. As noted by Lupton [316], with the advent of a ‘self-tracking culture’, “bodies, selves and data are entangled in the digitised self-tracking experience” (p. 5). Interestingly, such claims resonate with the views of philosophers and scientists who embrace an ‘extended’ (as in the sense of extended cognition) approach to issues of agency and the self [100, 103, 116, 232, 445].

A final point to note about the quantified-self movement is the emphasis that is placed on issues of self-enhancement and performance optimization. Issues of self-improvement are typically highlighted as one of the main reasons why individuals participate in self-tracking activities, and the quantified-self movement is often seen to provide new opportunities for self-related optimization. Swan [492], for example, suggests that:

Wearable electronics could facilitate automated self-tracking as vast amounts of data are uploaded to the cloud for processing by millions of agents (i.e., normal individuals) going about their daily lives. The QS [quantified-self] data ecosystem could include wearable electronics paired with mobile phones and cloud-based big data services so that the individual’s continuous personal information climate could provide real-time performance optimization suggestions. [492, p. 95]

Issues of self-improvement (in the context of the quantified-self movement) are typically approached from the perspective of the social sciences. Lupton [315], for example, notes that the emphasis on self-improvement tallies with neoliberalist narratives regarding the individual as a responsible citizen, willing and able to take care of their own self-interest and welfare. In addition to the social science issues, however, there are also a number of points of interest for both cognitive science and the philosophy of mind. In particular, the role of social and technological elements in supporting attempts at self-improvement raises important questions about the material character of agency, self-control and the will [100, 225, 378, 523]. In this respect, the sorts of feedback provided by (e.g.) digital dashboards might be seen to play

a crucial role in scaffolding our attempts at self-control, yielding implicit and explicit forms of behavioral guidance that we cannot help but heed. There is also a sense that in uploading information to the Web, and making it available for public viewing, we are effectively using the (online) social environment as a means to supplement, override, or perhaps even extend(!) the will [225, 378, 523]. As is noted by Lanzing [298], “being aware of the fact that *others* monitor one’s behavior adds another layer of externalized control and disciplining power” (pp. 10–11, original emphasis).

3.6 The Quantified Cognizer

In addition to issues of self-knowledge, the quantified-self movement also has implications for our understanding of human cognition. The notion of the quantified-self thus lays the foundation for what might be called *the quantified cognizer*: the idea that self-tracking activities provide us with an important opportunity to monitor aspects of human cognitive performance [289]. The value of such measures is typically discussed in relation to the ability to combine or merge data from multiple individuals. Consider, for example, the ability to monitor attentional focus using contemporary eye tracking devices [260, 291]. As noted by Kunze et al. [289], the use of such devices to derive aggregate measures of attentional focus (i.e., measures of collective attention) could be of practical significance in helping us understand the cognitive impact of socially-shared resources:

Content creators could... leverage quantified ‘mind logs’ to improve their works: Which part of a movie most excites viewers? What feelings does a particular paragraph in a book convey? At what point in a grant proposal does the reader lose interest? [289, p. 108]

This introduces us to an issue that will be discussed in more detail in Section 6.4. It concerns the use of the Web to glean information about the cognitive responses of multiple individuals. The general idea is that the nature of our engagements with the online environment provide insight

into our cognitive states and processes, and these can be combined to yield aggregate measures of collective or socio-cognitive processing. This idea is clearly related to the (rather more direct) use of the Web to further our scientific understanding of the human cognitive economy. The Web thus provides a valuable opportunity to gather data about the relationship between (e.g.) lifestyle factors (such as sleep patterns and dietary habits) and aspects of cognitive performance. Cognitive tracking apps, such as those marketed by Quantified Mind,⁵ provide us with one example of how the Web can be used to gather cognitive data. Another example is provided by computer gaming technology. Sea Hero Quest, for instance, is a mobile computer game that aims to gather information about the spatial cognitive abilities of both normal and clinical human subject populations [359]. The recent use of this game to assess the spatial abilities of over one million people [see 473] highlights the potential value of the Web as a platform for studying the human cognitive system.⁶

⁵<http://www.quantified-mind.com/>

⁶This use of the Web is not limited to the realm of individual-level cognition; it can also be used to shed light on the forces and factors that govern the performance of teams, groups and other multi-agent collectives [see 327].

4

Extended Cognition

Extended cognition refers to a particular form of cognition in which the mechanisms responsible for cognitive phenomena are deemed to extend beyond the neural realm. Work in this area tends to go by a variety of names, such as locational externalism [547, 548], active externalism [108], vehicle externalism [252, 413], environmentalism [412], cognitive extension [102] and the extended mind [108]. What all of these locutions have in common, however, is a commitment to the idea that aspects of the external, extra-neural environment can, at certain times, play a constitutive role in the material realization of the mind (or aspects thereof) (see Section 2.3.4).

Issues of extended cognition and cognitive extension have been a prominent focus of attention for philosophy of mind and cognitive science [5, 102, 106, 340]. However, extended cognition is also a topic that is of considerable interest to Web Science. This is reflected in a wealth of work that has emerged in respect of Web-based forms of cognitive extension [113, 214, 215, 229, 314, 447, 453, 454, 460, 475, 545]. One of the reasons extended cognition is important to Web Science is because cognitive extension is seen to lie at the root of an important form of cognitive enhancement for our species [102]. This makes the

Web a critical focus of attention for those who are concerned with issues of cognitive performance and cognitive well-being [e.g., 81, 312, 448, 538]. Inasmuch as we accept the idea that the Web forms part of the supervenience base for human mental states and process, then it is clearly important that we understand the implications of the online environment for our species-specific cognitive capabilities (both now and in the future).

The current section focuses on a number of strands of work that have emerged in relation to extended cognition and the Web. This is by no means a complete survey; for the terrain in this part of the intellectual landscape covers a bewildering array of topics and disciplines [see 454]. The areas of research reviewed in this section are those that have proved to be some of the more popular targets of philosophical and scientific attention.

4.1 Web-Extended Minds

The most explicit application of active externalist (or extended cognitive) theorizing to the Web comes in the form of the *Web-extended mind hypothesis* [447]. This is the idea that “the technological and informational elements of the Web can (at least sometimes) serve as part of the mechanistic substrate that realizes human mental states and processes” [447, p. 447]. Inasmuch as we accept the core claim of extended cognition—that the physical machinery of the mind can extend beyond the bounds of the biological body—the Web-extended mind hypothesis looks to be largely innocuous. Indeed, all the Web-extended mind hypothesis aims to do is take the conventional notion of extended cognition and apply it to the technological context of the Web. The question, however, is whether the Web is able to support real-world cases of extended cognizing. Are Web-extended forms of cognition a practical possibility, or are they merely a topic of theoretical interest and philosophical speculation?

One way of making progress on this issue is to refer to criteria that have been developed to evaluate putative cases of extended cognition. Perhaps the most well-known of these criteria (although by no means the only ones) are those proposed by Clark and Chalmers [108]

as part of their seminal work on the extended mind. The criteria in question are informed by a classic thought experiment involving a mnemonically-impaired individual, called Otto, and a conventional, paper-based notebook. In presenting the thought experiment, Clark and Chalmers [108] ask us to imagine two individuals: Inga and Otto, both of whom are situated in New York City. Inga is a normal human agent with all the usual cognitive competencies, but Otto suffers from a mild form of dementia and is thus impaired when it comes to certain acts of information storage and recall. To attenuate the impact of his impairment on his daily behavior, Otto relies on a conventional notebook which he uses to store important pieces of information. Otto is so reliant on the notebook and so accustomed to using it that he carries the notebook with him wherever he goes and accesses the notebook fluently and automatically whenever he needs to do so. Having thus set the stage, Clark and Chalmers [108] ask us to imagine a case where both Otto and Inga wish to visit the Museum of Modern Art to see a particular exhibition. Inga thinks for a moment, recalls that the museum is on 53rd street and then walks to the museum. It is clear that in making this episode of behavior intelligible (or psychologically transparent) to us Inga must have desired to enter the museum, and it is clear that she walked to 53rd street because she believed that that was where the museum was located. Obviously, Inga did not believe that the museum was on 53rd street in an occurrent sense (i.e., she has not spent her entire life consciously thinking about the museum's location); rather, she entertained the belief in a dispositional sense. Inga's belief, like perhaps many of her beliefs, was sitting in memory, waiting to be accessed as and when needed.

Now consider the case of Otto. Otto hears about the exhibition, decides to visit the museum, and then consults his notebook to retrieve the museum's location. The notebook says the museum is on 53rd street, and so that is where Otto goes. Now, in accounting for Otto's actions we conclude, pretty much as we did for Inga, that Otto desired to go to the museum and that he walked to 53rd street because that is where he believed the museum was located. Obviously, Otto did not believe that the museum was on 53rd street in an occurrent sense (Otto has

not spent much of his life constantly looking at the particular page in his notebook containing museum-related facts); rather, he entertained the belief in a dispositional sense. Otto's belief, like perhaps many of his beliefs, was sitting in the notebook, waiting to be accessed as and when needed.

Clark and Chalmers [108] argue that the case of Otto establishes the case for a form of externalism about Otto's states of dispositional believing. The notebook, they argue, plays a role that is functionally akin to the role played by Inga's onboard bio-memory. If this is indeed the case, then it seems to make sense to see the notebook as part of the material supervenience base for some of Otto's mental states, specifically his states of dispositional belief (such as those involving museum locations). The main point of the argument is to establish a (potential) role for external artefacts in constituting the physical machinery of at least some of our mental states and processes. If, as Clark and Chalmers [108] argue, the functional contribution of an external device is the same as that provided by some inner resource, then it seems unreasonable to restrict the material mechanisms of the mind to the inner, neural realm. It seems possible, at least in principle, for the human mind to occasionally extend beyond the head and into the external world.

As is noted by Clark [106], the aim of this particular thought experiment was to:

...convince the reader that, under certain conditions, the coarse functional role of a bio-external encoding could be sufficiently similar to that of a persisting internal encoding as to mandate similar treatment, revealing the non-biological resource as part of the physical machinery underpinning some of an agent's genuine mental states. (p. 448)

As Clark and Chalmers are quick to note, however, not every form of bio-technological or bio-artefactual coupling seems to invite treatment in cognitive terms. "There would," as Clark [96] suggests, "be little value in an analysis that credited me with knowing all the facts in the

Encyclopaedia Britannica just because I paid the monthly installments and found space for it in my garage” (p. 217).

What, then, are the conditions under which we should count a set of bio-external resources (such as Web resources) as proper parts of the machinery of the mind? In answering this question, Clark and Chalmers [108] proposed a set of criteria that have come to be known as the trust-and-glue criteria [see 105]. The criteria, as presented by Clark [105, p. 46], are as follows:

- “That the resource be reliably available and typically invoked.” [**Availability Criterion**]
- “That any information thus retrieved [from the resource] be more or less automatically endorsed. It should not usually be subject to critical scrutiny (unlike the opinions of other people, for example). It should be deemed about as trustworthy as something retrieved clearly from biological memory.” [**Trust Criterion**]
- “That information contained in the resource should be easily accessible as and when required.” [**Accessibility Criterion**]

One way of evaluating the extent to which the Web is able to support extended cognition is thus by reflecting on the kinds of interaction that we currently have with the Web. We can then assess the extent to which these kinds of interactive engagement comply with the requirements of the trust-and-glue criteria. In the remainder of this section, we provide an overview of the philosophical literature as it relates to the availability and accessibility criteria; trust-related issues are covered in a separate section (see Section 4.2).

The first of the trust-and-glue criteria relates to the availability of a resource. According to this criterion, one of the things that makes a bio-external resource apt for cognitive incorporation is the fact that it is “reliably available” and “typically invoked.” This seems to suggest that issues of portability and mobility are relevant to issues of cognitive extension: the more portable something is, the better its candidacy for cognitive incorporation. In this respect, the general thrust of technology development seems to speak in favor of Web-based forms of cognitive

extension [see 453]. This is because recent developments in mobile computing provide ever-more convenient ways of accessing the online environment. The portability of a smartphone, for example, is probably not all that dissimilar to the notebook resource that features in the classic Otto thought experiment [see 108]. Wearable devices may boast of even greater levels of portability. Consider, for example, the portability of a contemporary smartwatch, or the portability of future devices, such as electronic skin devices [556] and smart fabrics [275]. Given that the portability of such devices may begin to approximate the ‘portability’ of biological body parts, it seems increasingly likely that human Web users will be able to enjoy reliable access to the online realm.

A potential problem with this upbeat assessment of the availability criterion is that it ignores the power demands of mobile computing devices. As Heersmink [228] usefully notes, “digital diaries embedded in one’s smart phone, tablet, or other electronic device, in one sense, provide less reliable access, because they are inaccessible without electricity. So one not only needs to remember to bring the device when needed, but also to charge it when the battery is empty” (p. 586). In response to this particular worry, it may be worth considering research efforts that deal with issues of power consumption. This includes research into energy harvesting techniques [383], inexact/approximate computing techniques [219], novel forms of integrated circuit design [342] and wireless power transfer capabilities [290]. Another line of enquiry relates to the use of so-called backscattering techniques to yield battery-free devices that are able to use existing radio waves as the basis for wireless communication [198, 311].

Power and portability thus do not seem to represent much of a problem when it comes to a consideration of the availability criterion. There are, however, a number of other issues that are raised by the availability criterion. Crucially, these issues only tend to come to light when we seek to evaluate the availability criterion in the context of debates about the Web-extended mind [454]. To help us gain some insight into the nature of these issues, consider the claim that when it comes to issues of availability, it is not so much the portability of a resource that counts, as the fact that a resource is simply available

whenever we need it to be so. This idea is captured by the notion of *dynamic reliability*, which is discussed by Wheeler [543]:

What really matters is a property in relation to which portability makes a positive enabling contribution, but which may be secured without portability. That property is somewhat difficult to specify precisely, but, roughly speaking, it amounts to a kind of dynamic reliability in which access to the externally located resource under consideration is, for the most part, smooth and stable just when, and for as long as, that resource is relevant to some aspect of our ongoing activity. The qualifier ‘dynamic’ here reflects the fact that...organism-centred hybrid systems... often persist only when, and as long as, they are contextually relevant, meaning that the external resources concerned need not be smoothly and stably accessible at other times.

What Wheeler is suggesting here is that the continuous availability of a resource need not be all that important. Instead, what matters is the availability of a resource in a particular task context. To help us understand this idea, imagine that your only means of accessing the Web is via a desktop computer. Such a device is clearly not particularly portable, or at least, it is not as portable as, say, a smartphone. But what if you only need to perform a particular kind of task when you are actually seated in front of the computer? Does it really matter that the computer is not particularly portable, given the fact that it *does* permit access to the online realm whenever you need it to do so?

Another issue concerns the meaning of the term “resource” in the availability criterion. Note that when we focus our attention on the Web (or Internet), we are forced to recognize a distinction between a physical resource (e.g., a notebook or a smartphone) and the information that is made available as a result of interacting with that resource. In the original Otto (notebook) case, this distinction is of nugatory significance, because the relevant information is stored within the notebook. As a result, wherever the notebook goes the information is sure to follow. This, however, is not the case with a smartphone or wearable computing

device. In this case, the technological device could be readily available (and thus meet the availability criterion), but it could still fail to provide access to the relevant body of online information in certain situations (as when the networking capabilities of the device are disabled).

What this shows is that issues of availability in a Web-extended mind context are perhaps a little more complicated than we might have thought. In particular, when it comes to issues of cognitive extension, what seems to count is not the availability of a physical resource, such as a notebook or a smartphone; instead what counts is the availability of the information that is made available by the resource. This, however, yields a couple of problems. The first is that the availability criterion is in danger of being redundant. This is because the accessibility criterion already presupposes the availability of the target body of information—as long as some body of online information is sufficiently accessible, why concern ourselves with availability-related issues? The second (not altogether unrelated) problem is that exactly the same body of online information could be accessed via different devices in different contexts. In view of this, why should the properties of a *particular device* really matter that much, given that it is the availability of the information that really counts? We might, for example, access a particular body of online information via a desktop computer in the home environment, and then access the same body of information using a smartphone while on the move. The fact that one device is portable and the other is not does not seem to be of particular interest or relevance here. Much the same could be said for the appeal to typical invocation that also features as part of the availability criterion (see above). Given that the availability of online information may be unaffected by the fact that different devices are used on an occasional basis (at different times and in different places), why should we commit ourselves to the idea that typical invocation is of particular interest or relevance when it comes to Web-based forms of cognitive extension?

Enough has probably been said about the availability criterion at this point. Time, then, to turn our attention to the accessibility criterion.

In an extended mind context, issues of accessibility have proved to be something of a divisive issue. Ludwig [314], for example, suggests

that we can think of Wikipedia as extending our body of beliefs and knowledge as a result of the access it provides to specific kinds of information. Smart [447, 453] disputes this. He suggests that typical forms of browser-based access to Wikipedia undermine the extent to which we can talk of Wikipedia as an extended mind resource. There is a sense, of course, in which Wikipedia (*qua* encyclopedic resource) *does* enhance the accessibility of information. Compared to a paper-based version of the *Encyclopedia Britannica*, for instance, there can be little doubt that Wikipedia does provide improved access to certain kinds of information. In spite of this, however, the average HTML-rendered Wikipedia article is a veritable cornucopia of information, and the task of extracting relevant information from the article is one that requires considerable time and effort. Such temporal and energetic costs seem to count against the idea that Wikipedia content is accessible, at least insofar as our concerns relate to the demands of the accessibility criterion.

It is at this point that we encounter an appeal to the representational formats countenanced by the Semantic Web [46, 430] and linked data [53, 54, 226] communities. Alongside the familiar approach to representing online content in the form of HTML documents, there has been a recent attempt to move towards a more data-centric vision of the Web—one which emphasizes the role of the Web in providing access to globally-distributed bodies of machine-readable information [53, 226, 553]. Crucially, the representational formats that feature as part of this ongoing effort are ones that promise to transform the nature of our informational (and perhaps cognitive) contact with the Web. Linked data resources are thus able to serve as a back-end repository for all-manner of data-driven apps and services, many of which can be tailored to provide easy access to specific bodies of information content. To help us see this, contrast the way in which information is represented in Wikipedia—a system primarily designed for human consumption—with the approach taken by DBpedia [17, 55]. What DBpedia aims to do is represent a subset of the information content of Wikipedia in a form that is much more amenable to machine-based processing. DBpedia thus relies on the use of representational frameworks, such as Resource

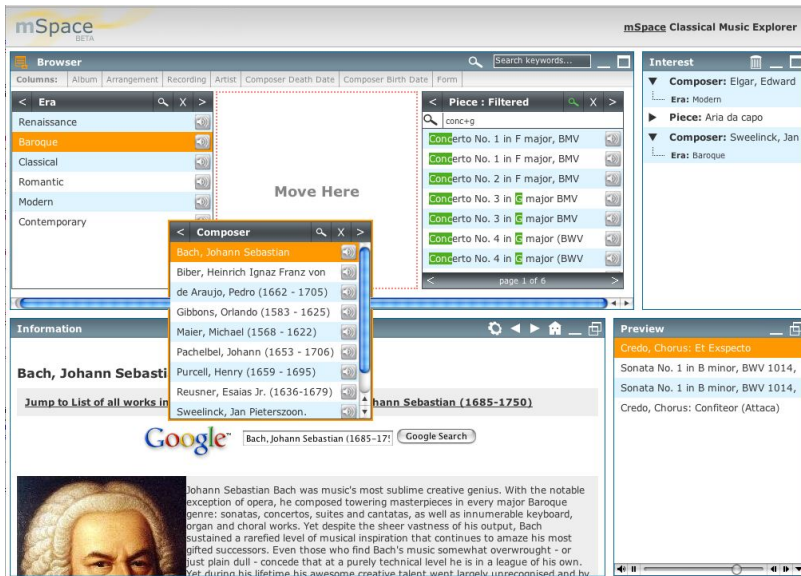


Figure 4.1: A screenshot of the mSpace Browser applied to the domain of classical music (source: <http://research.mspace.fm/projects/mospace>).

Description Framework (RDF), to create an interlinked network of information and data items, where the links represent semantically significant relationships between the data items. The advantage of this alternative representational scheme is that it effectively transforms Wikipedia into something of a structured ‘database’, improving the accessibility of specific items and supporting the flexible recombination of data items in particular task contexts. By itself, of course, this is unlikely to be of much use to a human agent: navigating a vast nexus of cryptically-encoded data items is hardly likely to be any easier than scrolling through endless reams of textual content on Wikipedia. But the benefits of DBpedia come into sharper focus once we consider the (many) ways in which the data can be used to support new forms of user interaction and information visualization [see 51].

One example of the benefits of data-centric representational formats is provided by schraefel et al. [424]. schraefel et al. describe a specific tool, called the mSpace Browser, which pulls together a variety of interaction and visualization techniques to provide rapid access to a wide

variety of heterogeneous sources in one persistent view (see Figure 4.1). Besides providing in-context information about any selected element in the domain, the mSpace Browser allows a user to manipulate the relationships that apply to elements within that domain, enabling them to easily reorganize and retrieve information as a means of addressing specific information requirements.

Data-centric representations not only provide us with novel ways of accessing online information, they also open up a wealth of opportunities to develop various forms of assistive intelligence. In order to help us see this, think about the way in which a body of semantically-enriched online data [see 430] might be exploited by a speech-enabled intelligent assistant, such as Apple's Siri or Microsoft's Cortana. Suppose that the body of data being exploited relates to the birth dates of significant historical figures. Given that the data is both machine-readable and annotated in ways that support various forms of (semantically-enabled) search and reasoning, it now becomes possible to initiate new kinds of epistemically-potent interaction with a conversationally-enabled digital companion. You can thus ask Siri questions such as: "When was Charles Darwin born?" Almost immediately, you will hear the response: "Charles Darwin was born the 12th February 1809." This mode of information access, it should be clear, is both easier and much more efficient than relying on browser-based access to Wikipedia.

The general idea, then, is that the move to data-centric representations opens up a range of opportunities for Web-based resources to be incorporated into cognitive processing routines [447, 453]. The present discussion also helps to highlight an important principle: when it comes to issues of cognitive extension, we should not restrict our attention to one particular kind of informational resource (e.g., conventional HTML Web pages). Different resources will, in general, yield profoundly different kinds of interactive relationship between the human cognizer and the online world, and this calls for a much more nuanced and refined view concerning the kinds of cognitive contact we have with the Web. In general, blanket statements that talk of the Web (in a general sense) as either supporting or not supporting extended cognition lack the sort

of precision that is required for them to be of genuine philosophical or cognitive scientific value.¹

4.2 Trusting The Web

Of all the trust-and-glue criteria, the trust criterion has proved to be the most problematic for the Web-extended mind hypothesis [403, 453, 454]. According to the original trust criterion, in order for something to count as a constituent element of an individual's cognitive economy, it should be "more or less automatically endorsed" and "deemed [to be] about as trustworthy as something retrieved clearly from biological memory" [105, p. 46]. This is often seen to present a problem for Web-based forms of cognitive extension, since the information that is retrieved from the Web is seldom information that we ourselves have curated. The result, it seems, is that we are seldom, if ever, presented with a situation in which online information can be trusted to the same extent as information retrieved from biological memory. It is for this reason that Clark [105] expresses doubt about the suitability of (at least some) Web-based resources for cognitive incorporation. Mobile access to Google, he suggests, does not count as a form of cognitive extension precisely because it violates the trust criterion.

The first thing to say about all of this is that it hazardous (to say the least) to make assumptions about the way Web users approach online information in the absence of any sort of empirical evidence. The fact that online information exists in a public and contested space [see 481] is not sufficient, by itself, to warrant conclusions regarding the extent to which we trust (or do not trust) online information. This is, in fact, something that is explicitly acknowledged by Clark [105]. The important question is thus not whether online information *ought* to be trusted (i.e., whether it is trustworthy); instead, the question is whether online information *is* trusted. This, it should be clear, is a question that can only be resolved on the basis of empirical research. In fact, a variety of strands of evidence suggest that Web users seldom

¹This also applies to work of a more general nature, such as work that seeks to understand the effect of the Web on cognitive performance [81, 333].

subject online information to critical scrutiny. Rather than engage in the meticulous evaluation of online information, Web users tend to rely on factors such as Web site design and navigability in forming credibility judgements [178, 179, 344]. Other kinds of cues, such as the number of followers a user has on Twitter, also tend to influence credibility judgements [541].

A second point to note about the trust criterion is that even if it could be shown that Web users do distrust the majority of online information, it is far from clear that the Web is congenitally unable to meet the demands of the trust criterion. Trust-related worries tend to be grounded in the fact that we do not have sufficient control over Web-accessible information. In other words, the concern is that online information is typically contributed by other users and is (perhaps) subject to communal editing efforts. But not all online information needs to be of this sort. Personal data stores [520] and private cloud computing repositories [202] illustrate some of the ways in which the Web may be used to establish something of a personal information ecosystem—one that provides access to carefully curated bodies of previously endorsed information. It is far from clear that such information is at risk of violating the trust criterion, at least to the same extent as information available on the ‘open’ Web.

Following on from the previous point, it is surely possible to imagine situations in which Web-based content was presented to a user in such a way as to avoid concerns about active scrutiny and automatic endorsement (the core concerns targeted by the trust criterion). Information could, for example, be presented to a user in a subliminal manner, thereby avoiding the route through conscious awareness [140]. Of particular interest, in this respect, is work relating to subliminal perception in human–machine interaction contexts [360]. As is noted by Negri et al. [360], subliminal stimuli can be used in concrete scenarios to bolster performance on mnemonic, problem solving, and navigational tasks.

The third, and final, point about the trust criterion relates to its validity. Should we accept the degree to which the trust criterion holds sway over our attempts to adjudicate the authenticity of putative cases of cognitive extension? One reason for scepticism, in this respect, concerns

the extent to which issues of automatic endorsement apply to the realm of inner (i.e., brain-based) information flows. When it comes to memory, for example, Michaelian [346] questions the extent to which mnemonic information is immune to evaluative assessment:

... due to the constructive character of encoding, consolidation, and retrieval, [mnemonic] records are not endorsed automatically upon retrieval—metamemory processes rather intervene to determine [the] endorsement/rejection of retrieved records. . . [346, p. 1156].

Michaelian uses this to cast doubt on the extent to which bio-external information can be regarded as a form of extended memory. However, the absence of automatic endorsement in the case of biologically-based memory could also be seen as undermining the legitimacy of the trust criterion. For why should we insist on automatic endorsement in the case of world-involving information flows, when such forms of endorsement do not appear to jeopardize the cognitive status of information flows in the inner, neural realm? Just to reinforce this particular point, it may help to reflect on situations where recall occurs in high-stakes situations. Imagine, for example, that you need to retrieve information from memory in order to make a critical decision—one where human lives are at stake. Even if you recalled the information with ease and believed the information to be true, would you not be inclined to check the veracity of the information by engaging in a form of ‘active scrutiny’, checking the information with respect to other (internal and external) sources of information? Does the fact that you are understandably circumspect about the retrieved information, in this situation, undermine the cognitive status of the processes that ‘served up’ the information? The answer to this is surely a resounding *no*. Given that we do not doubt the cognitive status of our biological memory simply because its informational deliverances are, in certain situations, subject to critical scrutiny, perhaps we should resist the urge to accord a significance role to active scrutiny and conscious evaluation in determining whether some bio-external resource counts as a *bona fide* element of an individual’s cognitive machinery.

Setting aside the specific details of the trust criterion, it is worth noting that issues of trust have a more general relevance to debates about the cognitive and epistemic impacts of the Web. One example of this comes from the epistemological literature. In particular, the extent to which information on the Web is of sufficient quality as to promote an overall elevation in our epistemic standing is an issue that has attracted the attention of the epistemological community [56, 314, 403]. Much of the debate centers on the extent to which the consumers of online content can trust the information they retrieve from the online realm. Record and Miller [403], for instance, note that it is incumbent on the users of smartphone devices to scrutinize the information they receive from those devices. They proclaim that smartphone users “implicitly trust their devices,” but they question the extent to which this trust is warranted. There is surely something right about this. In a “post-truth” era dominated by “false facts” and “fake news,” it seems that blind faith in the veracity of online information is unlikely to be a recipe for epistemic success. The result, according to Record and Miller, is that we are duty-bound to be suspicious of online content. Inasmuch as we accept that automatic endorsement is a precondition for cognitive extension, then it seems that the Web-extended agent is obliged to abstain from actions that are necessary to ensure the veritistic integrity of their doxastic system. The upshot is that in order to satisfy the trust criterion, Web users are required to forsake their duties as responsible epistemic agents.

It is not our aim, in the context of the present review, to contest the views of Record and Miller [403]; it is, however, worth noting that the trustworthiness of online information is a common focus of interest for the philosophy of mind, epistemology and Web Science communities. Indeed trust is one of the core topics that lies at the heart of the Web Science research agenda [188], and considerable effort is being invested in the development of techniques that can be used to assess the veracity of online information and improve its overall trustworthiness. It is here, perhaps, that we might be inclined to question the claims of Record, Miller and others. In particular, it is not always clear that the reliability of information is *always* undermined as a result of its being

located in a public and contested space. Indeed, there are occasions when the activities of others can help to indicate the reliability of online information [see 494]. Smart and Shadbolt [465] thus talk of the socio-technical construction of reliability indicators that can be used to guide epistemic evaluations [see also 39]. Something similar emerges in relation to work by Westerman et al. [541]. They show how system-generated cues (e.g., information about the number of Twitter followers associated with a user account) can be used to guide credibility judgements. Whether such credibility judgements are valid or not is, of course, another matter; but what matters for present purposes is simply the idea that communal access to information may provide opportunities for the socially-mediated epistemic evaluation of information content. Rather than seeing the social nature of the Web as a threat to positive epistemic standing, it may be that such properties help to assist the responsible, truth-seeking agent with regard to the epistemic vetting of Web-based information content.

As a means of further pressing this point about the epistemic utility of public access to socially-shared information, it may help to draw a parallel between online information and the information provided by road traffic signs. We uncritically endorse the information provided by road signs because the information is, in general, true. But part of the reason why these road signs are reliable is surely down to the fact that lots of people depend on them. It is, in other words, issues of social dependence that determine the continued reliability of the road signs. This does not, of course, mean that road signs are immune to vandalism. But given the role that road signs play in coordinating social behavior, we can at least be sure that considerable attention will be devoted to monitoring their integrity. At the very least, the chaos ensuing from any errors should suffice to arouse suspicion that something has gone awry, and thus lead to some form of remedial action. When it comes to the realm of road signs, therefore, we can see the necessary epistemic safeguards as, in some sense, *built-in* features of the larger socio-technical system. As more and more people come to rely on a particular information resource, so there is a greater need to ensure that any errors in the resource are quickly resolved. In these

cases, the resources are reliable precisely *because* they exist in a socially-shared space, not in spite of it. The general idea, then, is that far from undermining the trustworthiness of online information, public spaces can sometimes play an important role in ensuring the reliability of the information that is contained within them. The persuasive force of this argument is, of course, tied to issues of social penetration and social dependence. In particular, it is important that some information resource plays a critical role in ensuring the proper functioning of a social system. It might be thought that this serves to limit the applicability of the foregoing argument to the Web. However, the Web is a system that plays an increasingly important role in society, serving as a form of critical infrastructure for a broad array of social processes. As our dependence on this system increases, is there any reason to think that the reliability of (at least some parts) of the online realm will not match those of other resources that play a crucial role in ensuring the continued functioning of various forms of social machinery?

4.3 Dimensions of Cognitive Integration

In addition to the trust-and-glue criteria, there is another way of evaluating the Web-extended mind hypothesis. This approach comes in the form of the multidimensional framework proposed by Heersmink [228]. Heersmink suggests that cognitive incorporation—the integration of bio-external resources into an agent’s cognitive processing routines—can be evaluated with respect to a number of dimensions (see Table 4.1). In a Web-based context, these dimensions relate to the nature of the information flow between a human agent and some online resource, the accessibility of online information, the durability of human–Web couplings, the amount of trust a user puts into online information, the ease with which online information can be interpreted, and so on. According to Heersmink, “[t]he higher a situated system scores on the proposed dimensions, the more functional integration occurs, and the more tightly coupled the system is” [228, p. 577]. In essence, when it comes to Web-based forms of cognitive extension, we can use Heersmink’s multidimensional framework to evaluate the extent to which some form of agent–Web coupling should be seen as a genuine case of extended

Table 4.1: Dimensions of integration in Heersmink's [228] multidimensional approach to the evaluation of extended cognitive systems.

Dimension	Description
Information Flow	The dynamics of the information flow between an agent and a resource: one-way, two-way or reciprocal.
Reliability	The extent to which a resource is likely to be available.
Durability	The duration of the relationship between agent and resource.
Trust	The extent to which a resource is trusted by the agent.
Procedural Transparency	The ease with which a resource is exploited, often indicated by the degree of automaticity associated with the use of a resource.
Informational Transparency	The ease with which an agent can interpret and understand the information delivered by a resource.
Individualization	The extent to which the properties of a resource are tailored to suit the cognitive requirements of a given individual.
Transformation	The extent to which the representation and computational capabilities of the human cognitive system are transformed as a result of informational contact with the resource.

cognizing. In this case, the higher the rating that is achieved on each of the dimensions, the denser the integration between agent and online information, and thus the more likely that the online information is of constitutive relevance to an agent's cognitive states and processes.

Although the multidimensional framework has been applied to a number of computing technologies, including smartphone devices [248], it has not been used to examine the nature of our interactive engagements with the contemporary Web. As a means of addressing this shortcoming, we can apply the framework to the resources that were discussed in relation to the accessibility criterion (see above). These include a paper-based encyclopedia, Wikipedia, DBpedia, and Siri. Figure 4.2 shows the results of an informal analysis that applies the multidimensional framework to these resources (or at least the kinds of interactive engagement that are typically supported by these resources).

As is clear from Figure 4.2, each resource is associated with a distinct distribution of scores across each of the dimensions of the multidimensional framework. As a means of quantifying these differences, we can calculate what is called the Cognitive Incorporation Quotient (CIQ) for each resource. The CIQ, in this case, is a measure of the relative degree of cognitive integration/incorporation for each resource, which is calculated as the average of the normalized scores for each of the dimensions in the multidimensional framework. Based on the results of the present analysis (see Figure 4.2), it appears that DBpedia and Siri have the highest CIQs. This suggests that these resources provide relatively greater opportunities for cognitive incorporation, as compared to the case of browser-based access to Wikipedia and conventional access to a paper-based encyclopedia.

This is, of course, an informal analysis whose primary purpose is to illustrate the use of the multidimensional framework. It is thus not intended to support definitive conclusions regarding the 'extended' status of specific resources. In spite of this, the analysis helps to reinforce an earlier point about the dangers of blanket statements regarding the suitability of the Web as a target for cognitive incorporation. Consider, for example, that in using Wikipedia, Siri, and a DBpedia-coupled mSpace browser, we may be accessing the same (or at least very similar)

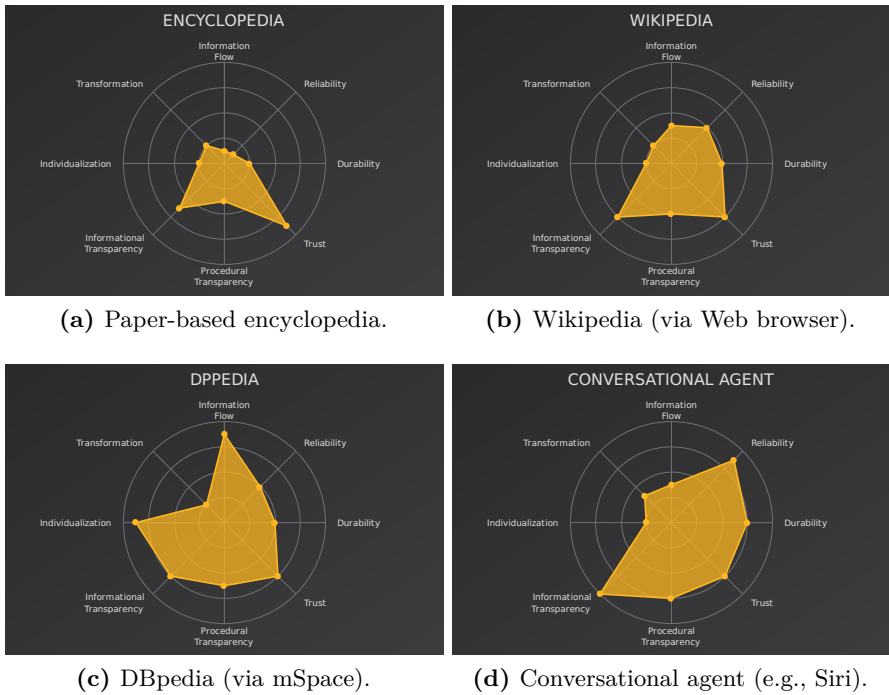


Figure 4.2: A comparison of different information resources using the multidimensional framework proposed by Heersmink [228]. [CIQ scores: Encyclopedia: 0.36; Wikipedia: 0.48; DBpedia: 0.64; Conversational Agent: 0.64]

bodies of information. The facts that serve as the target of interest for the extended mind theorist (e.g., Charles Darwin’s date of birth) may thus be contained in (or at least originate from) a common data source (e.g., the database associated with the Wikipedia system). What seems crucial, therefore, is the mode of access we have with some Web-accessible body of information, not the actual source of the information. This is clearly something that will depend on the nature of the devices and applications that are used to interact with the Web, and it is thus difficult (if not impossible) to state that some specific body of online information is apt for cognitive incorporation. It is for this reason that we caution against the use of blanket statements to the effect that the Web is (or is not) suitably poised to support cognitive extension. This

issue is not one that can be answered with any degree of philosophical precision in the absence of a detailed understanding of the specific ways in which human users interact with the online realm. Additional caution is required when it comes to the analysis of specific systems. In the case of Wikipedia, for example, it might be tempting to conclude that Wikipedia is unable to support cognitive extension on account of (e.g.) the accessibility of specific items of information (see above). The problem with this conclusion is that it is likely to be based on a particular way of accessing Wikipedia (no doubt one involving the use of Web browser technology). Given that whatever it is that is picked out by the term “Wikipedia” is distinct from our ways of accessing Wikipedia content, we are unlikely to be in a position to make general claims about the status of Wikipedia as an online resource that is (or is not) apt to support cognitive extension.

4.4 Extended Memory

The field of memory has become the crucible where many issues discussed in this section have been thrashed out or have reached their highest point of theoretical development. Many authors have discussed memory as a sort of proxy for the broader arguments about the validity and application of the extended mind [105, 418, 491], and in the form of E-Memory (electronic or digital memory technologies), it has been widely discussed as one of the main ways that the Web may extend the mind.

One reason for this focus is that new memory technology appears to be ubiquitous, and it is rapidly colonizing the cognitive niche largely vacated by paper-based media. From explicit external memory applications such as Evernote, to the general trend toward high capacity cameras and voice recording applications embedded in smartphones and tablets, to the still nascent trend toward wearable devices such as Google Glass, the world is becoming permeated with a host of technologies, all of which can encode, store and increasingly allow us to dynamically retrieve and organize memory traces.

Another reason is that mnemonic technology has a long history [148] and already is deeply incorporated in human cognitive ecologies. Knots in string, clay tablets, scrolls, codices, libraries and digital databases:

all can act as what Donald called exograms, i.e., artefactually based memory stores external to the body and brain (internally stored memory traces are what he calls engrams). The development of the human mind can be seen to march in step with the development of such technologies and may be directly implied in the development of new cognitive capacities [208, 321, 525]. Capstone technologies such as writing hold in place much of our civilization and give shape to the sorts of minds we possess and the sorts of beings we are [374], although these can and do undergo radical reshaping [112].

Donald [148] notes that one deep historical tendency in mnemonic technology is to externalize at least part of the memory processes from those flexible, labile but relatively interpenetrative ones natural to our brains, toward those that make use of the more durable and manipulable affordances of external objects, from clay tablets, to the printing press, to digital records. Donald argued it is through these extended memory technologies that our specifically human forms of memory can arise. But crucially, due to its historical nature, human memory is not fixed but can be reshaped and transformed as we create, appropriate and incorporate new regimes of mnemonic technology. Sutton's *complementarity principle* affirms that in extended cognitive systems, "external states and processes need not mimic or replicate the formats, dynamics or functions of inner states or processes" [491, p. 194]. Instead, we tend to deeply incorporate cognitive technology precisely when it makes available cognitive properties not found in our existing array of biological and technical memory systems.

There appear to be some deep continuities to the properties of external memory technology we develop but also, it is increasingly becoming clear, some important discontinuities. With the Web, we are passing through a major transition at least with respect to the technological basis of human memory. Whether this signifies a major transition for our cognitive processes and subjectivity more widely depends both upon the potential cognitive properties of this new technology, but also on how deeply we integrate this technology into our cognitive routines, and ultimately our minds.

Clowes [110] argues that four factors—*totality*, *incorporability*, *social entanglement* and *autonomy*—conjointly characterize much of what is new about E-Memory technologies. The factor which arguably is most novel with respect to the Web is *totality* [37, 110]. In part, totality is a question of capacity. To take one example, very recently a family might have one polaroid camera with which to record a lifetime's worth of memories within a single photograph album. Now, the same number of photographs may be taken in a day by a teenager with a smartphone. In principle, all of these photos—perhaps a lifetime's worth—could be uploaded to the Web and then retrieved and incorporated as needed in that individual's ongoing cognitive life.

In part, totality continues the historical tendency of memory technologies, which has been to make increasingly detailed, indelible and easily accessible memory traces that complement our more labile and reconstructive bio-memory systems. However, there are also important discontinuities. With traditional memory technologies, it has always required effort or work to lay down new records and, indeed, one generally had to make decisions about what to record. Thanks, however, to the convenience, voluminous capacity and sometimes automaticity of existing memory technologies, it is rapidly becoming possible to record any event we choose or just effortlessly and indiscriminately record in digital high-fidelity any memory trace that is available [37, 325, 334]. The tendency was first expressed in the trend toward lifelogging [222, 370], whereby users of digital memory technologies sought to record digital memory traces with the aspiration of capturing the totality of an individual's experience. This aspiration, conjoined with today's hyper-mobile and hyper-connected E-Memory gadgetry, means that a form of lifelogging is becoming an everyday reality for a substantial proportion of the population.

With the advent of cheap and efficient mobile network technologies, we can all record any thought, sound or image more or less as we choose. Moreover, the ability to organize and access digital memory traces using technologies such as Google Search or Amazon Alexa is rapidly changing the ecological space of memory. It is often simply easier to Google a disputed fact or personal memory trace than it is to

wrack our own brains [113]. As the user interfaces of mobile devices and their apps come to incorporate AI, personalization, and can be heavily customized to individual users' needs, our devices seem to be tending to meet the conditions on cognitive integration (as just discussed around Heersmink's [228] matrix) to an ever-increasing degree. As we habituate to, customize and individualize these devices and their set of apps, the patterns of information flow between us and them become ever denser, we often rely upon and trust their deliverances to a greater degree, and we can become more skillful interpreters of our preferred devices. Indeed, their usage becomes a form of second nature. In the process, the way that we accomplish a variety of tasks, such as navigating around a city, remembering a birthday or appointment, or keeping track of an exercise regime, is deeply transformed. This readiness by our mobile and wearable devices to be incorporated in a variety of cognitive operations is what Clowes calls *incorporability*. The new mobile devices are ripe for incorporation. The ecological space of memory and a raft of other cognitive operations is thereby being radically transformed, and with it, the potential to augment, or at least seriously alter the nature of human memory.

Some have worried that these deep tendencies have worrying implications for the human mind [81, 206, 515]. It may be that the ongoing tendency to deeply integrate E-Memory technologies may be undermining core human cognitive faculties. The thought is that as we come to rely on these technologies, our biological capacities will be weakened in the process. One of the more developed versions of this argument proceeds from the idea of transactional memory. Transactional memory is a form of memory organization found to occur in social groups where one participant will come to be seen as the expert in a particular field and then become treated as the memory specialist for other members of the group [536]. The process seems to be highly advanced in families and especially married couples where one partner might become the specialist on the names of films and who acted in them, and another might specialize in remembering to pay the bills [537]. In transactive memory systems a specialist takes over the role of remembering a specific type of information and is consulted by other members of the group when

information in that domain needs to be recalled. These spontaneous specializations seem to be a feature of the human social division of labor in respect of mnemonic tasks.

Wegner and Ward see the Internet as a supernormal stimulus that builds upon these tendencies in potentially dangerous ways [530, 531, 538]. Wegner and Ward argue that our brains tend to treat the Internet as just another sort of transactive memory partner, effectively an expert in everything, and in so doing, the danger is that our brains cease to store information internally and simply treat the Internet as a sort of universal expert. Indeed, there is some evidence that, at least when sitting at a non-Internet connected computer screen, we do tend to forget the details of some accessed information, but remember instead how to access it [472]. In this way, the thought goes, our memories and a wide range of our cognitive abilities may be undermined as we learn to rely on the sort of memory store provided by the Web, instead of the detailed information we might otherwise remember with our biological memory systems. The worry is this might be progressively undermining our organic memory systems. Ward adduces empirical evidence that, in the process, we are becoming self-deceivers with an over-enhanced sense of cognitive self-esteem, attributing knowledge to ourselves which we do not really know, but is merely accessible via the Web.

Clowes points out that today's hyper-mobile Internet users' self-attributions "might just be recognizing the new epistemic conditions as realities" [114]. That is to say, if per the extended mind argument, the deeply incorporated resources of the Web are ever-present, or at least dynamically reliable, then perhaps avid Web users are simply attributing to themselves knowledge that they can, under normal situations reliably access. Perhaps they implicitly take themselves to be Web-extended people. If so, we may have reasons not to worry so deeply about the new tendency to store more information in exograms and the expense of engrams. It might be better to epistemically evaluate the whole cognitive agent and not merely the agent's biological components. This could accord with the long history of the human mind which, as we have seen, tends to appropriate and incorporate a wide variety of artefactual systems.

Yet, it is perhaps hard to see why the vast knowledge stores of the Web should be considered as part of individual Web user's minds, even if all of the extended mind conditions are met. In part, this is because it is not clear that Web resources are best understood as personal resources at all. It might, some have argued, be better to consider some or all such resources as part of a cognitive commons that we can all draw upon, rather than as part of any individual's mind [113, 152, 483]. The Web may, at least in many scenarios, be more like a tube map to the London Underground, utilized in cognitive operations of many, but not part of any individual's mind, than they are like Otto's notebook: uniquely and personally part of Otto's cognitive operations.

Many of these arguments apply specifically to semantic memory. They become more challenging when we consider the way that E-Memory technology can support episodic and autobiographical memory functions in ways we have already touched on in the consideration of lifelogging [109, 232]. Such self-supporting—or self-instantiating—technology has already been shown to help some individuals keep a grip on themselves, and organize their lives in the context of traumatic brain trauma or conditions such as mild Alzheimer's [49, 153]. Experimental research with the SenseCam device, has shown that keeping track of one's everyday life by wearing a technology that automatically and unobtrusively takes a picture every few seconds, can help consolidate the memory of some patients [49] as well as enhance that individual's quality of life [437]. These potential benefits of E-Memory technologies, could be ported to more widely available portable and wearable devices given current and near-future Web-based technologies with the possibility of helping many people with a range of memory disorders.

But should such memory stores actually count as an instantiating part of anyone's self, or as part of them as a person? On a narrative theory of the self, the sorts of deeply detailed and deeply integrated personal memory stores that lifelogging makes possible may mean that the properties of some information stores and some Web-applications are better regarded as parts of an individual person or self than part of a mere technological adjunct [232, 445]. This may have serious ethical consequences, in for example, situations where an individual's data store

is compromised, when service is interrupted, or when memory stores are tampered with [230].

Some properties of the Web when considered as a memory store—in its current form—may countervail against it really being best considered as an adequate instantiation of a part of any individual’s self. Apart from the lack of dynamic reliability of much current Web-based mnemonic technology, E-Memory on the Web is often socially-entangled and autonomous in ways that may not be compatible with it being considered part of anyone’s self. When Facebook encourages someone to recollect a day or friendship and post it to their timeline, the prompt, is constructed by a series of hidden algorithms that largely depend on how others interact with that individual’s “personal” memory traces. The chance of recollecting any particular memory trace mediated by Facebook is based on massively socially-entangled memory systems [110]. Apart from this social entanglement and thanks to increasing permeation of AI technologies in the substrata of the Web, the Web is an ever more active and autonomous memory store.

Consider a personal photo, or video clip composed automatically from a series of photos, which unrequested pops up as you open your tablet device. The processes by which this particular photo is cued to appear to you, is dependent upon, rapidly changing, algorithmically complex and—at least to the majority of users—cognitively impenetrable processes which may, moreover, be commercially sensitive and thus hidden from view. From our foregoing discussion, it is contestable whether this should matter from the standpoint of the extended mind. They do not appear to figure either in the trust-and-glue conditions or in Heersmink’s multidimensional approach. The brain itself is cognitively penetrable only to a limited extent and what it recollects can also turn on processes that are socially driven. The algorithms cuing the memory trace are also likely to be highly personalized to data arising from that individual user and her immediate social circle. Yet, Web-mediated mnemonic technology appears to be ever-more autonomous and based on AI technologies driven by market forces, corporate policy and collective dynamics beyond the individual’s purview. Are such factors really irrelevant to whether a given device, application or process can be

counted as part of an individual? Arguably these considerations make it much more difficult to see contemporary E-Memory technologies as parts of individual memory stores at all, and more like a highly partial and civilization-specific forms of collective memory. A memory store may play a cognitively potent role for the individual, with important ethical and political consequences, without it needing to be considered part of anyone's mind.

4.5 Extended Knowers

The application of active externalist theorizing to the Web has potentially profound implications for our status as epistemic agents. In particular, arguments for the extended mind seem to suggest the possibility of extended knowers, i.e., agent's whose epistemic capabilities are not constrained by the limits associated with biological memory. In the case of an extended mind, an agent's beliefs are seen to be constituted (at least in part) by information that lies outwith the borders of skin and skull. All that matters, in this case, is the kind of access the agent has to external information. Given the right sort of access, external information should be counted as part of our body of beliefs about the world, and assuming these beliefs are true, it seems as though the Web-extended cognizer is poised to undergo a profound transformation of what they know about the world [see 314].

In order to help clarify all of this, it will be useful to consider a particular thought experiment. Imagine that you are equipped with a Web-enabled head-mounted display device (similar to the technological target envisioned by the Google Project Glass initiative). Such a device, let us suppose, comes equipped with a built-in camera and it is able to post images to a Web service and display results from the service directly within your field of view. Now imagine that the Web service is able to analyze photos of historical paintings and return information about the painting (e.g., the title of the painting and the name of the relevant artist).² By virtue of this technological set-up, you are

²See work by Johnson et al. [263] and Saleh and Elgammal [420] for details about the sort of technical approaches that might be relevant to this capability.

able to wander around an art gallery and view information about the paintings on display.

What should we make of your knowledge of art in such a situation? Would it be appropriate to conclude that you pretty much ‘know’ everything there is to know about the pieces of art contained in the art gallery, at least in terms of the information that is made available by the Web? Do you qualify as a *bona fide* art expert, despite the fact that quite a lot of the relevant processing seems to be occurring beyond the borders of the biological domain? Perhaps, as is often suggested in cases of extended cognition, the processing loop that travels through your technological devices and into the online environment should be considered as functionally akin to the brain-based neural circuits that would otherwise enable you to retrieve information from bio-memory. Inasmuch as this is the case, why should we seek to cast aspersions on your artistic expertise? If online information is poised to influence your thoughts and actions in roughly the same sort of way as information retrieved from bio-memory, then what is the basis for claiming that cases of brain-based information retrieval are indicative of a particular kind of cognitive competence, whereas the retrieval capabilities of some bio-technologically integrated ensemble fail to hit the (cognitive) mark?

For those who are still in some doubt as to the persuasive appeal of this argument, it may help to consider what it is that determines what you think you already know. What seems to determine whether we know or do not know something is not the fact that we are continuously, consciously aware of a set of relevant facts and figures. Instead, what seems to count is the kind of access we have to relevant information—the fact that when we need to recall information, it is there, easily made available to us by our bio-memory systems. But need our bodies of personal knowledge be so reliant on biologically-based modes of information storage? What if our access to externally-located information was the same as that afforded by bio-memory? In this case, it seems, there is no principled reason to suggest that the external information would not count as part of our own body of knowledge and beliefs about the world. As Clark [98] argues:

...it sometimes makes both social and scientific sense to think of your individual knowledge as quite simply whatever body of information and understanding is at your fingertips; whatever body of information and understanding is right there, cheaply and easily available, as and when needed. (p. 42)

In addition to whether it makes social and scientific sense to credit you with particular bodies of knowledge, we can also ask to what extent the mode of access outlined above would prompt a shift in your subjective feelings of what you did and did not know. If you were accustomed to having reliable, continuous access to particular bodies of information, would you eventually start to feel as though the externally-located information was actually part of the body of knowledge that you called your own? If someone standing next to you in the art gallery asked you whether you knew the identity of the artist for a particular painting, would you feel inclined to answer in the affirmative, based on your past experience of accessing art-relevant information?

One reason to expect a positive answer, in this case, stems from work relating to tool-related shifts in sensorimotor processing. In particular, one of the recent findings to come out of research on tool use is that changes to an individual's body schema can lead to changes in sensorimotor processing that persist beyond the period of time in which the tool is actually used [77, 78] (see Section 3.3). Cardinali et al. [77], for example, provide evidence to suggest that the use of a 40cm long mechanical grabber introduces a change in the kinematics of subsequent grasping movements. They report that the use of the grabber alters the kinematics of subsequent free-hand grasping movements and pointing movements in a way that suggests the represented length of the arm has been modified. Such results point to something of an 'after-effect' or 'capability echo' concerning the impact that a previously integrated tool has on our action-related capabilities.³

Intuitively, we might expect the cognitive equivalent of a capability echo to assume the form of a subjective shift concerning the nature of

³Further evidence in support of this 'after-effect' or 'capability echo' is found in work using virtual reality applications [52, 209, 476].

our cognitive and epistemic capabilities. In other words, we might expect cognitive incorporation to alter our sense as to what we feel we can do, even when we are decoupled from the bio-external resource that grounds our extended cognitive abilities. This is a view expressed by those who claim that research on bodily incorporation serves as the corporeal equivalent to the notion of cognitive incorporation [99]. Clark [100], for example, claims that reliable access to external bodies of information is all it takes to effect a shift in capability-related subjective experiences:

Easy access to specific bodies of information, as and when such access is normally required, is all it takes for us to factor such knowledge in as part of the bundle of skills and abilities that we take for granted in our day to day life. And it is this bundle of ‘taken-for-granted’ skills, knowledge, and abilities that. . . quite properly structures and informs our sense of who we are, what we know, and what we can do. [100, p. 106]

Such claims are of particular interest in light of recent findings concerning the effect that Web access has on our subjective sense of our cognitive and epistemic capabilities. In situations where people use the Web to search for online information it appears that their subjective sense of what they know is transformed, extending to include not just the information that is available from bio-memory but also the information that is available from the online realm [173, 530, 531]. Searching for information online, Fisher et al. [173] suggests, thus “leads people to conflate information that can be found online with knowledge in the head” (p. 675). Similarly, Ward [530] notes that as people turn to the “cloud mind of the Internet, they seem to lose sight of where their own minds end and the mind of the Internet begins. They become one with the cloud, believing that they themselves are spectacularly adept at thinking about, remembering, and locating information” (p. 88). Interestingly, and in accord with the notion of a capability echo, such effects persist beyond the period of time in which users have access to the Internet. Thus, even when a mobile phone is not at hand, the mere fact of being accustomed to its use renders an agent overconfident in estimating the limits of her epistemological capabilities [530, 531].

Such results tend to speak in favor of the conclusion that access to the Web alters the nature of the metacognitive judgments underpinning our subjective sense of what we do and do not know. Not all studies, however, support this conclusion. Ferguson et al. [169], for example, report the results of a study in which access to the Internet appears to undermine rather than enhance feelings of knowing.

The available research thus presents a somewhat confusing picture regarding the effect of Web access on meta-memorial processes. Nevertheless, it seems clear that access to the Web can effect a fundamental shift in subjective experiences regarding epistemic capabilities. In future research, it will be important to relate the findings of psychological experiments with theoretical models that detail how metacognitive processes alter as the result of past experience [161, 408]. It will also be important to pay close attention to the kinds of access that users have with the online environment. To echo the point made above (see Section 4.1), we should not assume that the results obtained with one kind of Web access (e.g., that provided by Google Search) will generalize to other kinds of Web access. The cognitive impacts of the Web are likely to depend on the modes of access we have to the Web, and much the same may be true of the metacognitive processes that give rise to a rich array of epistemic feelings and emotions [14, 347].

4.6 Supersized Knowers?

The idea that the Web might, on occasion, give rise to extended knowers has led epistemologists to make a number of claims regarding the impact of the Web on our epistemic profile. One implication of the Web-extended mind concept, Ludwig [314] argues, is that we are able to envisage a profound transformation of our doxastic potential. In particular, he anticipates “an explosion of dispositional beliefs and knowledge that is caused by digital information resources such as Wikipedia or Google” (p. 355). Similar views are expressed by Bjerring and Pedersen [56]. They argue that the Web enables us to enjoy various forms of “restricted epistemic omniscience,” corresponding to more-or-less “complete knowledge about a particular, fairly specific subject matter” (p. 25). We thus arrive at a claim that seems to follow quite naturally from the

possibility of Web-extended minds: cognitively-potent forms of Web-based bio-technological bonding will lead to states of Web-extended knowledge, in which many of our epistemically-relevant doxastic states supervene on material elements that are typically treated as part of the technological and informational fabric of the Web.

Unfortunately, there are a number of problems confronting this sort of claim. One of the most pressing problems relates to the way in which attempts to satisfy the criteria for cognitive extension (e.g., those proposed by Clark and Chalmers) seem to cause problems for the epistemic integrity of the technologically-extended agent [453]. Consider, for example, the accessibility criterion, which was discussed in Section 4.1. The general idea behind the accessibility criterion is that external information should be quickly and easily accessible—it should be possible for agents to access information and incorporate it into their problem-solving routines with little effort or difficulty. Accessibility thus appears to demand a degree of *fluency* with respect to the kinds of interaction an agent has with bio-external resources, where the notion of fluency can be understood (at least in part) as the “subjective experience of ease or difficulty with which we are able to process information” [375, p. 237].

Now, the problem with claims regarding easy access and fluent interaction is that these properties seem to be in some tension with the possibility of Web-extended forms of knowledge. Thus, one of the findings to emerge from research into fluency is that fluent processing leads to a truth bias, in which the ‘truth’ of some body of external information is judged relative to the subjective ease with which it is processed [10]. Fluency thus seems to speak in favor of the possibility of Web-extended minds, but it seems to work against the interests of Web-extended knowers (i.e., agents whose epistemic credentials are enhanced as a result of Web-based forms of cognitive extension). We thus encounter what has been dubbed the *extended cognizer vs. extended knower problem* [453]:

Extended Cognizer vs. Extended Knower Problem

The properties that work to ensure that an external resource can be treated as a candidate for cognitive incorporation are also, at least in some cases, the very same properties

that work to undermine or endanger the positive epistemic standing of a technologically-extended agent.

This is an important problem; for it suggests that Web-extended minds may not be the natural harbingers of Web-extended knowledge. Although we might be inclined to assume that, as Clark [107] puts it, “extended minds breed extended knowers” (p. 3760), things seem to be somewhat more complicated in view of the extended cognizer vs. extended knower problem. Indeed, if the proposed tension turns out to be real, then this will undoubtedly undermine claims about the potential role of the Web in effecting a significant transformation of our epistemic capabilities [56, 314].

The extended cognizer vs. extended knower problem makes its presence felt in a number of ways. In fact, the problem has already been encountered in relation to the earlier discussion on trust (see Section 4.2). Recall that as part of that discussion, we encountered the idea that the trust criterion imposes a constraint on extended cognition, such that the extended cognizer should eschew the deliberate and meticulous evaluation of online information [403]. The upshot, claim Record and Miller [403], is that the extended cognizer is (by definition) unable to engage in the sort of epistemic vetting and active scrutiny that, they insist, is mandated by the questionable reliability of online information. In the absence of such scrutiny (which, recall, is a condition for cognitive extension), the beliefs of the Web-extended agent seem to lack the necessary justificatory status that would otherwise warrant claims of positive epistemic standing. In other words, by virtue of the specific properties of online information, the extended beliefs hosted by the Web-extended mind are unlikely to count as a form of extended knowledge.

A similar sort of worry occurs in relation to claims regarding the use of linked data formats to enhance the accessibility of online information. Recall that in Section 4.1 we encountered the idea that linked data formats provide a means to address the demands of the accessibility criterion [447, 453]. The problem with this idea is that it potentially undermines the role of ‘peripheral’ information in guiding credibility evaluations. By peripheral information, we mean the

kinds of cues and affordances that are associated with the context in which information appears. Such information includes, for example, the visual design elements of a website, the structural organization of information, and the presence of socially-constructed reliability indicators and system-generated cues (see Section 4.2). Crucially, all these kinds of informational cues, prompts are affordances are used by human users in forming epistemically-relevant judgments about online content [179, 344, 465, 541]. The concern, therefore, is that if we remove this epistemically-crucial contextual information as a result of the adoption of data-oriented representations, we will remove an important source of information that undermines the extent to which an individual's beliefs can be said to be justified, and thus indicative of positive epistemic standing. In an attempt to enhance the conditions under which Web-extended minds emerge (i.e., using data-centric representations to enhance the accessibility of online information), we may inadvertently undermine the extent to which the Web-extended cognizer can be regarded as a Web-extended knower.

Here is a somewhat more subtle manifestation of the extended cognizer vs. extended knower problem. Suppose we take the findings of Ward [530, 531] and Fisher et al. [173] at face value and accept that the Web inflates our sense of what we know. If this is indeed the case, then it seems that the Web-extended agent is likely to regard themselves as more knowledgeable and expert when they have access to the Web. Now, it might seem that such subjective shifts are epistemologically innocent and do not affect our overall epistemic standing. However, according to a recent view in social psychology, situations that alter expertise-related self-perceptions may contribute to shifts in cognitive style. In particular, according to the *earned dogmatism hypothesis*, individuals who believe themselves to be more expert in a given domain are also more likely to adopt a dogmatic, closed-minded orientation [377]. This is important, because dogmatism is often seen to undermine the epistemic standing of an individual. This is especially so in the case of virtue-theoretic conceptions of knowledge within contemporary epistemology [31, 204]. According to a position known as virtue responsibilism, for example, knowledge ascriptions are tied

to the possession of a number of positive character traits or intellectual virtues [21, 558]. Some notable examples of such virtues include intellectual patience, open-mindedness, intellectual humility, curiosity, intellectual courage, epistemic diligence, intellectual tenacity, carefulness and intellectual conscientiousness. Perhaps unsurprisingly, dogmatism and closed-mindedness are not to be found among this list of intellectual virtues. Instead, dogmatism and closed-mindedness are deemed to constitute intellectual vices, and these are almost universally denigrated on account of their knowledge-undermining potential.⁴ In particular, dogmatism seems incompatible with positive epistemic standing, since it seems that the beliefs of the dogmatic individual are unlikely to track the truth. Even in situations where the dogmatic individual does succeed in believing the truth, the truth of such beliefs seems to be due more to luck than it does anything else. The true beliefs of the dogmatic individual will thus violate an epistemic anti-luck condition [392, 394], and this is almost universally regarded as inimical to knowledge attribution.

It seems, therefore, that subjective shifts in our sense of what we do and do not know are of considerable epistemological importance. Far from being epistemologically innocent, the subjective changes that accompany our exposure to the Web may modify elements of our cognitive character, leading us to fall prey to intellectual vices that work to degrade and diminish our status as epistemic agents.

4.7 Privacy Implications of the Web-Extended Mind

Given that all forms of cognitive extension involve resources that lie beyond the traditional biological borders of the skin and skull, it should be clear that extended cognitive processes are of greater visibility to external agencies than are their brain-based counterparts. Although

⁴As always, there are dissenting voices. Smart [456], for example, suggests that while we are correct to see dogmatism as a genuine form of intellectual vice, its vice-like properties are only apparent at the individual level of analysis—the level of individual agents. In contrast, when we turn our attention to the collective level, the dogmatic properties of individual agents have a more virtuous feel to them. This is due to what Smart refers to as *mandevillian intelligence*—the idea that individual cognitive shortcomings, limitations or biases can, on occasion, play a positive functional role in yielding collective forms of cognitive success (see Section 7.7).

future forms of technological innovation may support a form of ‘brain reading’, whereby we are able to monitor and interpret patterns of brain activity [267, 507], it seems as though extended cognition permits a particularly easy form of observational access to an individual’s cognitive machinery. In particular, cognitive extension seems to afford a more-or-less ‘direct’ observational route to an individual’s cognitive ‘innards’. By extending the machinery of the mind to the extra-organismic realm, it becomes possible to access and monitor a subset of the physical elements that realize an individual’s mental states and processes—to quite literally ‘see the mind in action’ [see 286].

The accessibility of bio-external cognitive circuits leads to a natural concern with covert forms of ‘cognitive veillance’.⁵ The idea, here, is that cognitive extension opens the door to forms of monitoring and surveillance that are not encountered in the case of brain-based forms of cognition (i.e., neurocognition). Even if we allow for the possibility of future brain-reading technologies, it seems unlikely that such technologies will introduce the same sort of privacy concerns as those associated with extended cognition. This is because, in the brain reading case, it seems much harder (although not impossible) to imagine situations in which a particular individual was being monitored without their knowledge. In the extended mind case, however, covert monitoring could be as easy as looking over someone’s shoulder as they perform a long multiplication task using pen and paper resources.⁶ The visibility and accessibility of extended cognitive routines thus ‘opens up’ the mind to a particularly worrisome form of what we might call ‘cognitive privacy violation’.

Unfortunately, such worries are only amplified in the case of Web-extended cognition; for issues of privacy and surveillance are of particular importance when it comes to the nature of our relationship with the Web [369]. Crucially, the Web provides a wealth of opportunities to record and monitor the actions of individuals as they engage with the online realm. This is important, because the Web is, to a large extent, a public space where bodies of data and information are easily accessed

⁵This extends the types of ‘digital veillance’ identified by social scientists [see 317, chap. 2].

⁶See Wilson and Clark [550] for an active externalist approach to long multiplication.

and shared, sometimes without our knowledge. Every time we interact with the Web—to post a message or perform a Web search—we face the risk of making some body of personal information available to a variety of interested parties, e.g., government agencies, commercial organizations, and (in the worst case) criminal fraternities. It should now be clear why Web-based forms of cognitive extension are likely to accentuate concerns about cognitive privacy: once we accept the idea that cognitive extension provides a means to observe the mind in action, then the Web emerges as the perfect arena in which covert forms of cognitive veillance could be undertaken, perhaps on an industrial scale. There is, of course, no reason to assume that such forms of surveillance will be benign in nature. Indeed, they could be undertaken for a variety of highly nefarious purposes. We thus come face-to-face with a vision of a surveillance society that is, in some respects, just as unsettling as the one portrayed by Orwell's [376] classic novel, *Nineteen Eighty-Four*. The citizens of Orwell's society were at least able to enjoy some degree of privacy concerning their innermost thoughts and feelings. Such is not the case when the machinery of the mind extends to include the global digital networks of the online world.

The notion of the Web-extended mind thus raises important issues concerning cognitively-relevant forms of surveillance, monitoring and control. Importantly, such issues *only* arise when we seek to situate the notion of cognitive extension within the specific socio-technical context of the Web. In fact, in the absence of the Web, the notion of extended cognition might seem to be relatively benign from a privacy perspective. Similarly, in the absence of an extended mind perspective, the existing debate over privacy and surveillance lacks an important, and politically-significant, cognitive dimension. Given the recent interest in various forms of surveillance that have been undertaken by government agencies [see 251, 469], now is arguably a good time to consider the privacy implications of the Web-extended mind.

5

Embedded Cognition

Unlike the notion of extended cognition, which sees the elements of the extra-organismic environment as sometimes playing a constitutive role in the realization of cognitive states and processes, the notion of embedded cognition rejects the idea that the boundaries of cognition extend beyond the traditional biological borders of the human agent. This is not to say that proponents of embedded cognition fail to recognize the cognitive significance and relevance of the extra-organismic environment. As is noted by a leading advocate of the embedded approach:

... cognitive processes depend very heavily, in hitherto unexpected ways, on organismically external props and devices and on the structure of the external environment in which cognition takes place. [417, p. 393]

The distinction between embedded and extended cognition thus relates to the constitutive relevance of resources that lie beyond the biological borders of the individual. The extended theorist accepts the idea that bio-external resources can sometimes qualify as part of the machinery of the human mind, while the embedded theorist insists that cognitive mechanisms are resolutely confined to the intra-cranial realm,

even though they accept that such mechanisms may be subject to causal influence by aspects of the extra-organismic environment.

Given that embedded cognition recognizes that human cognitive processes are subject to causal influence by extra-organismic factors, it should come as no surprise to learn that there are many of points of interest for the embedded cognition theorist when it comes to the Web. In fact, practically all the features of the Web that surface in relation to embodied and extended cognition also come into focus when we view the Web through the lens of embedded cognition. There are, however, a number of specific points of interest for the embedded cognition theorist. Some of these are reviewed in the current section. Section 5.1 presents the notion of the *cognitive commons*, which is a term that has been used to describe the Web from an embedded perspective. Subsequent sections cover the Web of Things (WoT) (see Section 5.2), the notion of programmable worlds (see Section 5.3), and the cognitive implications of the Web from an embedded perspective (see Section 5.4). Finally, in Section 5.5, we discuss why the conceptual distinction between extended and embedded cognition is important from a scientific, philosophical and engineering perspective.

5.1 The Cognitive Commons

In Section 4, it was suggested that the informational and technological elements of the Web could serve as part of the physical substrate that realizes cognitive states and processes. Needless to say, not everyone is convinced by this proposal. Harnad and Dror [152, 221], for example, explicitly reject the notion of cognitive extension. They see cognition as something that is confined solely to the inner, neural realm, and they dispute the idea that cognition extends beyond the metabolic surfaces of the biological agent. “[C]ognition,” insist Harnad and Dror [221], “begins and ends at the cognizer’s sensor and effector surfaces.”

Although Harnad and Dror reject the notion of extended cognition, they do appreciate the potential significance of the Web when it comes to issues of human cognizing. The Web, they suggest, corresponds to a publicly accessible informational environment whose contents are both socially-created and socially-maintained. In essence, Harnad and Dror

see the Web as a “cognitive commons”—a resource that is collaboratively created and poised to alter the cognitive power and potential of humanity at both an individual and collective level:

... the worldwide web, a distributed network of cognizers, digital databases and software agents, has effectively become our “Cognitive Commons,” in which distributed cognizers and cognitive technology can interoperate globally with a speed, scope and degree of interactivity that generate cognitive performance powers that would be inconceivable within the scope of individual local cognition alone. [152, p. 21]

The idea of the Web as a cognitive commons is an appealing one, not least because it emphasizes the social nature of the Web—the fact that the online environment emerges as the result of collective efforts. This is important in helping us understand some of the ways in which the Web can support human cognition. In exploiting the resources of the online environment, we are thus sometimes able to perform cognitive tasks that would have been difficult, if not impossible, to accomplish should we have attempted to perform those tasks in the absence of ‘social’ support. An example comes in the form of scientific open access initiatives, which allow scientific papers to be publicly accessible. Such initiatives, in conjunction with the use of global information networks, may help to stimulate various forms of creative insight and intellectual progress. As Harnad [220] rightly notes, the Web allows us to accomplish something akin to “scholarly skywriting”—scientific theories, thoughts, ideas, and sometimes research data, are made available in ways that are increasingly accessible to fellow academics and scientific colleagues. It is almost as if the outputs of scientific enquiry were written in the sky for all to see.

One virtue of this particular mode of information distribution and dissemination is that it enables thoughts and ideas to influence the minds of individuals with very different outlooks, knowledge and experience. As is noted by Clark [96], such forms of information flow potentially enable a community of individuals to establish shortcuts or conduits between thoughts and ideas that would otherwise be too distant or

disconnected to be linked by a single individual. In discussing the role of language as a means to support the migration of ideas between different individuals, Clark thus suggests that:

... such migrations may allow the communal construction of extremely delicate and difficult intellectual trajectories and progressions. An idea that only Joe's prior experience could make available, but that can flourish only in the intellectual niche currently provided by the brain of Mary, can now realize its full potential by journeying between Joe and Mary as and when required. The path to a good idea can now criss-cross individual learning histories so that one agent's local minimum becomes another's potent building block. [96, pp. 205–206]

When we merge this idea with the claims of Harnad and Dror, we are presented with a powerful and compelling vision concerning the cognitive potential of the Web. In particular, we can see the Web as boosting the scope and scale of the processes alluded to by Clark, supporting discontinuous jumps within a space of thoughts, theories, ideas, and insights.

5.2 The Real World Web

As first sight, the claim that we can understand the cognitive significance of the Web from the standpoint of embedded cognition looks to be largely uncontroversial. The appeal to the notion of 'embeddedness' is, after all, intended to draw our attention to the role of the extra-organismic environment in shaping the cognitive processing routines of a given individual. Inasmuch as the Web counts as an extra-organismic environment (which it surely does), then it seems reasonable to assume that the Web is relevant to issues of embedded cognition. Moreover, such an approach allows us to connect our understanding of Web-embedded cognition to a long history of theory and research that details the cognitive significance of tools, technology and other aspects of material culture [96, 321].

There is, however, a potential problem with all this. The problem relates to the distinctive nature of the online and offline environments—between the virtual environment of the Web and the physical environment of the real-world. The typical focus of embedded cognition concerns the nature of our interactive engagements with the physical environment, and there seems little doubt that we are embedded in such environments. But is the same true of the virtual environment? Are we embedded in the online world to the same extent as we are embedded in the offline one? This is, in fact, a difficult question, not least because it highlights a potential deficiency in our understanding of the notion of embeddedness. Inasmuch as we see the online environment as an environment that is distinct from the physical world—some abstract informational realm that lies beyond the browser interface—then there is always a risk that the Web will be seen to be of little relevance to embedded cognition.

We have, of course, already encountered one response to worries of this sort in the section on embodied cognition (see Section 3.1). In that case, the response centered on the validity of the distinction between the online and offline environments. Such a distinction, it was suggested, might be called into question as a result of the way in which Web-based information is used to guide everyday thoughts and actions.

Another kind of response is inspired by research into the WoT [210, 211, 212, 561]. For present purposes, we can see the WoT as the Web-based counterpart to the conventional Internet of Things (IoT) [207, 354]; the only notable difference being that the WoT relies on the use of Web-based communication protocols to effect the exchange of information. Aside from this (perhaps minor) technical detail, we can think of the WoT as providing similar functionality to the IoT. As with the IoT, for example, the WoT aims to extend the reach of the Web into the wider physical environment, equipping a variety of everyday material objects with data acquisition, data processing and data exchange capabilities. This is important, because it helps us see the Web as more than just some remote informational realm that is separate from our everyday engagements with the real-world. In place of this vision, we are encouraged to view aspects of the real,

physical environment—the world in which we are materially embedded—as themselves falling within the empirical and theoretical remit of contemporary Web Science. The result is that the putative relevance of the online/offline distinction, especially when it comes to issues of embedded cognition, starts to look increasingly spurious. For why should we downplay the relevance of the Web from an embedded perspective when it seems that an ever-expanding array of real-world physical objects will be Web-enabled? Such forms of enablement not only raise questions about the material nature of the Web environment (e.g., do Web-enabled objects exist as literal parts of the Web?), they also provide us with another way to think about the cognitive significance of the Web. This becomes apparent when we consider the way in which the WoT promises to transform the properties of everyday physical objects. As physical objects become invested with computational and communicative potential, so their implications for human thought and action are likely to shift. In essence, we can think of the WoT as altering the affordances of objects and thus changing the resultant topography of the “enactive landscapes” [see 276] in which our cognitive routines are situated.¹

As a concrete example of the way in which the WoT may help to shape human cognition, consider the case of prospective memory. Prospective memory is a form of memory that involves “remembering to carry out intended actions without being instructed to do so” [20, p. 343]. The case of an individual who needs to remember to defrost the meat by removing it from the freezer when they return home from work serves as a typical example of prospective memory. Such forms of memory are, of course, relatively commonplace, and they are pretty much indispensable in terms of our ability to coordinate our lives effectively—a fact that is all too sadly evidenced by those suffering from impairments in prospective memory [554]. As has been pointed out by a number of commentators [e.g., 475, pp. 36–37], the advent of smart

¹The notion of an enactive landscape refers to the way an agent perceives the world in specific task contexts [276]. The basic idea is that enactive landscapes consist of affordances that support the expression of particular actions. Such affordances change according to the goals of the agent; however, they also change in response to the use of tools or exposure to specific technologies.

environments provides a range of opportunities to radically reshape the nature of prospective memory. Staley [475], for example, talks of smart devices being used to implement prospective memory systems that allow individuals to “embed their intent” within specific environments, such as within their home or office. An individual could thus be reminded of the need to engage in particular actions (e.g., to remove the meat from the freezer) whenever they are suitably placed to perform those actions (e.g., when they first enter the kitchen upon returning home from work). Such interventions may be of particular importance when it comes to the ‘treatment’ of individuals with cognitive impairments, such as individuals suffering from dementia or some form of acquired brain injury. Smart objects thus provide us with an important opportunity to build on the existing use of the Web for the purposes of providing clinically-relevant cognitive support. When it comes to prospective memory, for example, McDonald et al. [337] found that the use of a specific Web technology (namely, Google Calendar) could serve as an effective memory aid for individuals suffering from impairments in prospective memory. The WoT enables us to extend the scope of such interventions to the local ambient environment that surrounds specific individuals. Indeed, the use of smart objects in a clinical setting may encourage us to engage in a productive re-evaluation of what it means to ‘treat’ some form of cognitive impairment. Instead of the usual emphasis on individual-centered interventions (e.g., the administration of drugs), we can also begin to note the value of changes to the local environment in which an individual is embedded. This is an idea that is perfectly compatible with a situated or ecological approach to cognition. As one proponent of the ecological movement once suggested, if you want to understand human intelligence you should ask “not what’s inside your head, but what your head’s inside of” [319, p. 43]. Such sentiments highlight the value of interventions that aim to enrich the environment, as opposed to interventions that aim to repair the brain: by making the world smart, we are able to trade ecological enrichment for neuro-computational sophistication, adapting the environment to suit the residual capabilities of the cognitively-impaired individual.

5.3 Programmable Worlds

From the perspective of embedded cognition, the WoT promises to transform the environment into something of a cooperative partner when it comes to the performance of particular tasks.² In this sense, the WoT speaks to the sort of vision that is often alluded to in discussions of ubiquitous computing [539] and ambient intelligence [535]. The general idea is that by extending the reach of the Web to the elements of our physical environment, we are able to transform the environment into something that is more responsive to our needs and more supportive of our thoughts and actions. This is, of course, the guiding principle behind work that seeks to develop a seemingly endless array of ‘smart things’, e.g., smart TVs, smart cars, smart buildings, smart cities, and smart environments [see 354].

How do we create devices that are able to adapt their modes of operation to suit the (cognitive) needs and requirements of particular individuals? One answer to this question is to support individuals in customizing the behavior of smart objects. Such approaches aim to give individuals greater control over their environments, enabling them to tailor the structure of Web-based information flows to suit their idiosyncratic needs and concerns. There is an obvious parallel, here, with issues of personalization and customization in the cognitive science literature. Sterelny [483], for example, talks about the role of personalization in smoothing the flow of cognitively-relevant informational exchanges between agent and artefact. Similarly, Clowes [113] discusses the importance of personalization to issues of cognitive incorporation. In this case, Clowes suggests that we should see personalization as enabling

²Castelfranchi et al. [84], for example, talk of the attempt to “create environments that are sensitive and responsive to inhabitants’ needs and capable of anticipating their needs and behavior as well” (p. 17). Something similar is proposed by Ricci et al. [406]. They suggest that the development of networking technologies, cloud technologies and the IoT supports the emergence of “a new kind of smart space in which digital, physical, and social layers are strongly intertwined. These spaces extend the classic assistive functionality of AmI [Ambient Intelligence] toward more proactive possibilities, where the smart environment not only monitors people as they perform tasks, or supports them by executing their requests, but also influences and changes their plans and intentions” (p. 60).

a technological resource to be ‘assimilated’ into a agent’s cognitive routines. Issues of personalization and customization also surface in relation to the notion of cognitive niche construction [102, 482, 483], which is defined as “the process by which animals build physical structures that transform problem spaces in ways that aid (or sometimes impede) thinking and reasoning about some target domain or domains” [102, p. 62]. Here, the ability to modify the environment is seen to play an important role in yielding certain kinds of cognitive success. By virtue of the ability to restructure the environment in particular ways, an agent is sometimes able to transform a difficult (or impossible) problem into something that is more suitability aligned with the native information processing repertoire of the bare biological brain.

One way in which individuals can customize the behavior of Web-enabled devices and influence the shape of Web-based information flows is exemplified by technologies that use rules to implement specific actions in response to specific informational contingencies. An illustrative example of this sort of technology is provided by IF This Then That (IFTTT).³ IFTTT is a Web service that can be used to connect a broad array of online systems (e.g., Gmail and Dropbox), social media sites (e.g., Twitter, Facebook, Instagram, YouTube and LinkedIn), and Internet-enabled devices (e.g., Fitbit and Phillips Hue⁴). These are collectively referred to as “services” by the IFTTT community. The connection between services is realized by trigger-action rules, which as their name suggests consist of both a trigger component and an action component. The trigger part of the rule is a conditional expression that defines a set of events and contingencies associated with a source service. When these conditions are met, the action part of the rule is ‘triggered’. The action part of a trigger-action rule engages with a target service to implement certain operations. For example, the action might use the Application Programming Interface (API) provided by a social media site to upload or modify content. IFTTT thus provides a means by which specific actions can be performed in response to changes in the global online data environment. Crucially, this environment includes the

³<https://ifttt.com/>

⁴<http://www2.meethue.com/en-gb/>

Table 5.1: Examples of IFTTT trigger-action rules. Note that these examples are operational applets that can be accessed via the IFTTT website (<https://ifttt.com/>).

#	Trigger (IF)	Action (THEN)
1	The International Space Station passes overhead	Change the color of my Phillips Hue light bulb
2	Rain is forecast for tomorrow	Send me an email
3	I take a photo with my smartphone	Upload the photo to my Flickr account
4	A photo is uploaded to my Flickr account	Post the same photo to my Facebook page
5	My Fitbit device registers low sleep levels	Remind me to go bed early
6	I receive a Google Scholar email alert	Append the contents of the email to a running note in Evernote

data streams that emanate from a broad array of physical devices, and this provides an opportunity for IFTTT users to detect and monitor changes that occur in both the physical and virtual environments. Other kinds of data streams can be exploited to indirectly detect events in the physical environment. For example, a user could create a rule that notifies them whenever a new article is added to the “Earthquake in California” category on Wikipedia.

Some concrete examples of IFTTT rules are listed in Table 5.1. These rules highlight a number of important features of IFTTT. The first thing to note is that IFTTT rules can exploit the data streams that are provided, managed and maintained by a large number of different agencies. The first rule (Rule #1) in Table 5.1, for example, exploits a data stream that is managed by the U.S. space agency. As is illustrated by this rule, the sorts of contingencies targeted by IFTTT need not always be restricted to the online realm: they can also refer to states-of-affairs that exist in the physical environment. Here we can see the value of the Web in providing a digital representation of the physical environment. By creating data streams that are sensitive to the changing

nature of specific parts of the physical environment, individuals are afforded the opportunity to maintain an awareness of events that, in many cases, would be beyond the reach of their perceptual systems. As is demonstrated by Rule #2, such an awareness can effect changes in behavior, enabling individuals to coordinate their thoughts and actions with current or future states of the world. There is, in addition, no reason why the focus of awareness should be limited to states-of-affairs that exist beyond the boundaries of the biological body. As is evidenced by Rule #5, IFTTT rules can be authored to take into account the information provided by Internet- or Web-enabled wearable devices. Not only do such devices transform the nature of our informational contact with our own bodies and behavior, they also afford an opportunity to connect our biological body with an ever-expanding suite of sensors, actuators and online computational services.

A second feature of IFTTT is revealed by Rules #3 and #4 in Table 5.1. This concerns the ability to support complex chains of rule execution. The action (or THEN) part of Rule #3 involves the uploading of a photo to Flickr. This, in turn, triggers the execution of Rule #4, which posts the uploaded photo to a Facebook page. This scenario involves the successive execution of just two rules; however, there is no limit to the number of rules that can be linked or ‘chained’ in this sort of way. Similar kinds of rule chaining can also be applied to the realm of smart objects and devices, thereby enabling an individual to alter the properties of the local (or remote) environment in response to some contingency that occurs in the physical or virtual world.

At this point, it should be clear that IFTTT rules enable the casual user to engage in a form of ‘programming’. This is perhaps most obvious in the case of rules that automate certain kinds of tasks and routines (e.g., Rules #3 and #4). Beyond this, however, we can see how the advent of smart objects enables the physical environment itself to be approached as a programmable resource. We thus encounter the notion of what might be called the *programmable world*—the idea that the Web and Internet enable the physical environment to be the target of forms of control and manipulation that are typically the preserve of more conventional programming efforts [see 493]. Superficially, there is

nothing particularly radical about this idea; for the kinds of control and manipulation that are being achieved in these cases are still very much grounded in standard forms of computational processing. A program that executes against the physical environment (via the use of smart objects) is thus similar to the kinds of programs that execute against a set of online digital resources—in both cases the relevant program exploits digital representations that are made available by the Web. Despite this, the programmable world concept remains important, if only because it highlights the opportunities that individuals have to configure the local environment in ways that are germane to their idiosyncratic needs, interests and concerns.

A similar point applies to the management of information flows that exist solely in the online realm. Rule #6, for example, demonstrates how information from one service (email) can be automatically routed to another service (Evernote), resulting in the aggregation of information in a single location. This is a simple example, but it is an example that highlights an important point: the use of IFTTT rules can be used to control the flow of information in the online realm, enabling an individual to assemble novel information resources and customize their cognitive contact with such resources. This extends a point that was encountered earlier in the review (see Section 4.1) regarding the extent to which conventional HTML pages are able to support the emergence of Web-extended minds. In the context of that discussion, we highlighted the danger of limiting our attention to one particular form of informational contact with the Web. The Web, we suggested, provides a range of opportunities for individuals to access particular bodies of information, and this potentially undermines the validity of blanket statements regarding (e.g.) the suitability of the Web as a target for cognitive incorporation. The present discussion of IFTTT reinforces this particular point. In particular, we are now in a position to see how the informational ecology of the online environment can be customized in ways that befit the cognitive requirements of a specific individual. An individual could, for example, create a rule that presents them with particular bodies of information using a simple text message. In this case, there is clearly no need for the individual to actively search

and retrieve information from the Web. Instead, the information is presented to the individual whenever they deem it appropriate for that information to be presented. There is, in addition, no need for us to limit our attention to the case of simple textual prompts. As is illustrated by Rule #1 in Table 5.1, bodies of online information can be used to effect changes in the local ambient environment, yielding perhaps subtle (or even subliminal) forms of cognitive and behavioral influence.

5.4 Embed with the Web?

The cognitive implications of being embedded in a world of Web- or Internet-enabled smart objects is rarely subject to detailed examination. Yet, it is clear that many human cognitive processes are liable to change in response to the constant availability of the Web and the plethora of objects that rely on the Web to support their functionality.

Technologies, tools and material culture, more generally, can sculpt, shape and support a range of cognitive abilities [321]. For some, this relationship is particularly intimate and aspects of material culture are seen to shape and constrain the development of the human brain. The result is that at least some of our cognitive abilities, along with their neural instantiations, might be bound to particular artefactual substrates, with tools and technology constricting our cognitive abilities and directing their development down predetermined paths.

One of the controversies to emerge from debates in this area is whether the material support provided by the paper book allows for the development of literary capabilities and whether these capabilities are hindered by the emergence of alternative media ecologies, such as the Web. It has been suggested that the new digital culture might undermine the brain circuits that support such abilities [551]. Indeed, an unintended consequence of the Web may be a progressive diminishment of the very cognitive abilities that allowed us to create the Web in the first place.

One motivation for this idea stems from Dehaene's hypothesis that new cognitive abilities can only be developed if the brain circuits that instantiate them are sufficiently similar to allow them to be "recycled" with respect to a new task context [138]. The idea is that as our brains adapt to a particular ecological niche (or, by extension, human-made

tools), certain circuits of the brain will need to be appropriated for new tasks, and in the process, reshaped. Dehaene's hypothesis is that the information processing properties of brain circuits will be altered by the appropriation and recycling to new cognitive tasks associated with skills built around novel tool use. However, the acquisition of artefact-dependent skills may come with a cost. In particular, as new skills are acquired, others may be lost, as the co-opted brain region no longer supports the previous task so well. Dehaene has put this forward as a zero-sum hypothesis: gaining one set of skills through neural recycling is in danger of undermining the skills that were previously supported by the relevant neural substrate. This leads to a worry about the cognitive impact of Web technologies. In particular, as our brains are 'reprogrammed' by our exposure to Web technologies, the fear is that we might become less able to acquire skills that were more easily forged in an earlier (pre-Web) era.

It is important to say here that despite finding that our brains can indeed be rapidly reshaped by the Internet [312], there is little or no evidence that this involves a correlative reduction in cognitive abilities. The zero-sum hypothesis mentioned by Dehaene is thus without empirical basis, at least when it comes to the use of Web technologies. Moreover, it is clear that the cognitive ecology of the Web is currently undergoing rapid transformation, at least in part to make it possible to fluidly develop skills such as reading [115]. The invention of E-Readers and tablet computers seems to provide a suitable medium for much traditional style reading. Indeed, recent empirical studies suggest that the cognitive effects of new screen devices are likely to be negligible, especially when it comes to reading-related abilities. Many of the effects originally reported for poor comprehension or slower retention in respect of Web-based reading are now much harder to find, with some studies reporting little or no difference between electronic devices and paper books [517].

It seems that some of the rhetoric surrounding the 'impact' of digital technologies on our cognitive abilities has all the hallmarks of a moral panic, which fails to consider the extent to which technologies change to support existing as well as newly minted skills. It remains an open

question as to just how much human cognitive abilities are tied to the details of our contemporary technosphere; however, there are good historical examples of complex skill sets making transitions between material substrates. The ability to read was arguably not seriously compromised as we moved from hand-copied codices to printed books.

It is also seldom considered which new cognitive capacities might become available as human beings develop skills to manipulate our Web-based material culture. As mobile and wearable technologies become ever-present props, the resources of the Web start to provide a sort of ever-present virtual environment for cognition. From the presence of virtual agents such as Amazon's Alexa—ready to answer questions and perform tasks for us—to the highly personalized apps resident on our mobile devices, a broad array of Web-based tools now provide a new sort of constant virtual cognitive environment. It is one that we can often interact with through natural language, which responds in highly tailored ways to our individual cognitive profiles—in part thanks to the huge data sets companies such as Google now build about all of us—and that can be used for an ever-expanding set of cognitive tasks.

5.5 Extended vs. Embedded: Does It Matter?

It might be thought that the preceding discussion of smart objects, smart environments, the cognitive commons, and so on, is relevant *only* to embedded cognition. This is, in fact, a mistake. The WoT, for example, could be deemed to be of as much relevance to extended cognition as it is to embedded cognition, especially when one considers the opportunities that Web-enabled objects provide for cognitively-potent forms of bio-technological bonding. Similarly, much of the content in Section 4 could be approached from an embedded rather than an extended perspective. It might thus be thought that the distinction between embedded and extended approaches to cognition is of little practical (or indeed theoretical) significance. Although the distinction between extended and embedded approaches has been the source of a lively philosophical debate [2, 104, 105, 417], why should members of the Web Science community (and, more generally, the computer science community) care about such a distinction?

In the current section, we aim to address this particular issue by introducing the reader to the five *E*'s: explanation, ethics, enhancement, education, and engineering. By considering the implications of extended and embedded cognition for each of these topics, we hope to show that the distinction between extended and embedded cognition is more than just a philosophical sideshow or an inconsequential terminological dispute. Table 5.2 provides a summary of the key points raised in respect of this issue.

5.5.1 Explanation

Inasmuch as we see human cognitive capabilities as rooted in the operation of extended cognitive mechanisms, then it should be clear that an extended approach to cognition will be required to help us understand the material underpinnings of the human cognitive system. In the absence of such an approach, there is a risk that human cognitive phenomena may appear somewhat mysterious, appearing to arise from (neural) mechanisms that lack the wherewithal to realize them.

Donald [148] refers to this sort of 'explanatory gap' in his discussion of human memory. The human cognitive system, Donald argues, has evolved to support the externalization of memory by means of written language, number systems, diagrams, and other representational systems. These external memory systems, he suggests, play functionally identical roles to those encountered in the case of biological memory. The result is that we cannot understand human mnemonic capabilities in the absence of an account that refers to the properties of external representational systems, as well as the nature of the interactions we have with these systems. An account that limits its attention to (e.g.) the processing dynamics of the biological brain, is, it seems, unlikely to yield a complete understanding of our cognitive capabilities, and it is for this reason that Donald and others encourage an approach that expands the unit of cognitive scientific analysis beyond the limits of the biological brain [254, 255, 256, 257].

A similar sort of worry sometimes surfaces in debates about the extended mind. Quite often, when explaining someone's actions, we refer to the beliefs and desires possessed by the agent. Such mental states

Table 5.2: The distinction between extended and embedded cognition is important when it comes to the five *E*'s: explanation, ethics, enhancement, education, and engineering.

Explanation	If human cognition is grounded in the operation of extended mechanisms, then an extended approach to cognition is required to understand human cognitive capabilities. An embedded perspective, in this case, risks leaving certain aspects of human cognition unexplained or, in the worst case, inexplicable.
Ethics	If human cognition is extended, then bio-external resources from part of the machinery of the human mind, just as do parts of the biological brain. From an extended perspective, therefore, the deliberate destruction of an external resource perhaps counts as a form of personal assault.
Enhancement	Cognitive extension may lead to qualitative shifts in cognitive processing that are not easily predicted or explained from the perspective of embedded cognition.
Education	An extended perspective has implications for educational policy, especially as it relates to the use of the Web in academic examinations.
Engineering	The distinction between extended and embedded cognition matters when it comes to the engineering of intelligent systems. If human cognition is extended, then a system that emulates the processing capabilities of the biological brain is unlikely to exhibit human-level intelligence.

typically form part of the (folk psychological) explanatory story that we use to make sense of an agent's behavior. If, however, some of our beliefs are stored in external media, then perhaps they too should count as part of the explanatory story. Extended mental states, suggests Chalmers, "can function in explanation in very much the same way that beliefs function, and they should be regarded as sharing a deep and important explanatory kind with beliefs. This explanatory unification is the real underlying point of the extended mind thesis" (p. xiv) (see foreword to Clark [102]). Chalmers thus argues that external beliefs are essential to explaining human thought and action. Inasmuch as we see the Web as forming part of the realization base for extended mental states, then it seems we may be obliged to refer to the Web as part of our attempts to develop folk psychological explanations of human behavior.

5.5.2 Ethics

Clark [98, chap. 7] points out several moral and societal issues related to cognitive technology. These include unequal access to technology; unauthorized access to personal information by individuals, corporations, and governments; and the potential for alienation. These issues also apply to Web-based forms of cognitive extension. We do not all have equal access to the Web, resulting in information 'haves' and 'have-nots' [242, 487]. Our Web-based extended memory systems might be accessed by others, resulting in an infringement of our privacy. If our external memories are publicly accessible, then we should perhaps be more careful in generating certain online content. Others, however, have pointed out that this may lead to self-censorship, which limits human freedom and autonomy [335]. Spending too much time online, the argument goes, may alienate us from friends, family, nature, and political engagement. Conversely, the Web may also foster friendship [65] and political engagement [63].

Heersmink [231] draws attention to three additional ethical issues related to extended cognition, namely, the consequences of cognitive technology for brains, cognition, and culture; the moral status of cognitive technology; and the implications of cognitive technology for issues of personal identity. Some theorists have argued that an overreliance on

online information diminishes some of our onboard cognitive capabilities and alters our cognitive profile in perhaps undesirable ways. Technology critic Nicholas Carr [81, 82], for example, argues that relying on the Internet as an external memory system reduces the amount of facts we store in long-term memory, thereby making us less knowledgeable. Susan Greenfield argues along similar lines. She writes: “Those with more facts at their immediate disposal can build richer constructs of reality and thus have a world view informed by a context that enables deeper understanding, more *wisdom*” [206, p. 221, original emphasis]. On her view, Internet use generates a culture in which people have a shallower understanding of topics and are less wise. Carr and Greenfield’s criticisms are based on research in cognitive psychology, which shows that when information is available externally, people tend to spend less effort committing such information to biological memory [472]. This is referred to as the “Google effect.” Heersmink [229], however, has questioned the generalizability of Sparrow et al.’s research to the Internet, as their research does not actually involve using the Internet, it only involves storing statements on a desktop computer. It is very likely that the Web transforms human biological memory systems in all kinds of ways, but the currently available empirical research in cognitive psychology does not seem to support strong negative conclusions of the kind voiced by Carr and Greenfield.

If the Web extends the human cognitive system, then parts of the Web are likely to have a particular moral status. On an extended cognition view, altering one’s extended memory is on a par with altering one’s hippocampus (the part of the brain involved in memory). Johnny Søraker claims that “the case with Otto’s notebook suggests that information and information technology can have moral status, but only if they are constitutive and irreplaceable in a strong sense” [471, p. 14] (see also [304]). The same reasoning applies to Web-based extended cognitive systems, implying that online information may have moral status. Therefore, we ought not interfere with people’s online minds. Carter and Palermos [83] argue that, for these reasons, stealing or otherwise intervening in one’s extended mind should be seen as a form of personal assault. An extended approach to cognition thus has ethical

and legal implications, and these often go beyond those associated with an embedded approach to cognition.

Web-based extended memory systems have an obvious relation to human selfhood. Some philosophers argue that our self is a narrative construct, arguing that who we are as persons is essentially our (unfolding) life story [423]. The building blocks of the narrative self are specific autobiographical memories. Web technology can play an essential role in storing such memories. Lifelogging, for example, aims to capture and store personal experiences in an indexed database called a lifelog. Gordon Bell's MyLifeBits project, developed at Microsoft, is an interesting example. Bell's digital lifelog contains photos, videos, Web pages, GPS locations, letters, memos, receipts, legal documents, business cards, meeting agendas, symposia programs, diplomas, employee evaluations, newspaper clippings, childhood drawings, birth certificates, and much more. He can search his lifelog relatively effectively and uses it as an external autobiographical memory. Bell and Gemmell predict that lifelogs "will become vital to our episodic memory. As you live your life, your personal devices will capture whatever you decide to record. Bio-memories fade, vanish, merge, and mutate with time, but your digital memories are unchanging" [37, p. 57]. Because of the supplementary role of Web-based lifelogs for autobiographical memory, they also play a key role in our narrative self. Other Web-based technologies that play an important role in our narrative self are social media systems such as Facebook, Twitter, Flickr, etc. The photos and status updates provided by these systems arguably play an important role in autobiographical memory [143], and thus contribute to the construction of the narrative self.

5.5.3 Enhancement

A consideration of cognitive extension is required to understand how to enhance human cognitive performance. In a careful reflection on the properties of extended cognitive systems, Wilson [549] writes:

One way for cognitive extension to lead to a reduction in functional capacity is through cognitive clutter: by adding

more bells and whistles to an existing cognitive system, we might well cause it to operate less effectively, or even lose certain kinds of functionality; so-called “smart technologies,” for all the benefits they bring, are often used in ways that have this effect for particular tasks (e.g., in driving). [549, p. 23]

We should thus be aware that not all forms of cognitive extension are cognitively beneficial. In some cases, they may be detrimental. There is so much information available on the Web that it does not always lead to the desired epistemic effect. Part of the reason is that we do not always have the cognitive resources available to filter and evaluate search results, Rich Site Summary (RSS) feeds, emails, and so no. Information overload can thus lead to the use of false or incorrect information. We should be aware of such effects of technology on our cognitive capacities. Generally, however, using technology is cognitively beneficial [365]. As Clive Thompson puts it: “From Socrates onward, we lose old cognitive skills as we gain new ones, but we’ve benefited from the trade-off” [501, p. 135].

Having access to online information has transformed many aspects of human cognitive life, usually for the better, including education, navigation, journalism, and academic scholarship. We can perform significantly more cognitive tasks when we have access to the Web. We look up information with search engines, store documents in the cloud, navigate with online maps, read online newspapers and books, engage with online courses, use online recipes, check online timetables, watch online videos, and so on. The Web allows us to streamline many of our daily cognitive tasks. The human brain is a powerful information processor, but there are limits to its capacity to remember, calculate, navigate, plan, and learn. To overcome these limits, we create and use artefacts, including the Web, and these enable us to perform cognitive feats that we might otherwise struggle to achieve.

As we have seen in our discussion of the complementarity principle (see Section 4.4), human beings tend to rely on tools that provide capacities that contrast with their existing cognitive abilities. Insofar as the Web provides such tools, we can expect them to be increasingly

factored into everyday cognitive skills. From the extended cognition viewpoint such skills can, in principle, be viewed as capacities of the agent; whereas on the embedded viewpoint, these appear more like the agent outsourcing his or her capacities. The deep incorporation of tools can look like the agent effectively swapping out “real” internal cognitive capabilities for non-cognitive external ones. In this way, adopting extended or embedded viewpoints can have real implications for issues of cognitive enhancement and diminishment [113]. Choosing to adopt a theoretical viewpoint can also have ethical consequences, in that the adoption of the very same technological device may look like a form of diminishment from an embedded viewpoint and a form of enhancement from an extended viewpoint.

5.5.4 Education

Clark writes that “[t]echnological education will be crucial if human-machine cooperation is to enrich and humanize rather than restrict and alienate” [98, p. 183]. To optimize our embodied interactions with cognitive technology, we should thus educate pupils and students so as to enable them to use technology in the best possible way. Theorizing on the role of technology in education from an extended mind perspective, Mike Wheeler [543] writes:

Perhaps what we ought to focus on, then, is the education of those hybrid assemblages, a focus which is entirely consistent with the goal of endowing the brain with the skills it needs to be an effective contributor to such assemblages. From this perspective, of course, there are extremely good reasons to support the increased presence of technology in the examination hall.

Wheeler thus argues that we should focus on educating extended cognitive systems and allow students to use technology when they are taking exams. If we couple with online information to form extended cognitive systems, how should such systems be educated? We should first of all be aware that, in some cases, extending one’s cognitive system may have undesirable effects. Consider, again, Wilson’s [549]

observation that, in some cases, extending one's cognitive system results in performing tasks less effectively, or may even result in losing certain kinds of functionality. Web-users should be educated so as to prevent or minimize such negative consequences. Moreover, as pointed out above, one's extended memory such as, for example, a Web-based application such as Google Calendar, can be accessed by others, potentially resulting in an infringement of one's privacy. It is important to educate people as to how to use Web-based technologies in the best possible way and to make people aware that their Web-extended minds might be accessed by others [404].

Epistemologist Duncan Pritchard writes: "Focussing on the real-world cognitive situations that citizens encounter—situations which are these days laden with technology—is entirely the right approach for our educational policies to take" [395, p. 50]. Given that the Web is ubiquitously used in many contexts, we should focus on educating Web-extended cognitive systems. This means we should teach pupils and students information literacy skills such as the ability to efficiently navigate, evaluate, compare, and synthesize online information [145, 229, 405]. Given that search engines are the main portal to the Web, it is important to educate people how to use search engines in an epistemically responsible way by being able to define search queries, choose the best search results, and evaluate sources for reliability and validity [see 145].

5.5.5 Engineering

The distinction between extended and embedded cognition alters the way we think about the development of Web-based systems. The majority of development efforts are, of course, geared towards supporting human agents with respect to the performance of cognitive tasks [145, 364, 365], and this highlights the relevance of cognitive ergonomics [241] and cognitive engineering [57] to the development of Web-based systems. This is the case irrespective of whether we view cognition from an extended or embedded perspective. On an extended view, however, the engineering of online systems provides the material fabric for extended cognitive mechanisms, and it is in this sense that technological and software engineering can be seen as, quite literally, a form of 'cognitive

engineering’ or ‘mind engineering’—a way of providing the material elements from which extended cognitive organizations are built. Such a perspective obviously affects our approach to the engineering of technological systems. As is noted by Blomberg [57], an extended approach to cognition can “help researchers and system designers to achieve a shift of perspective that could have practical consequences downstream from research and design” (p. 98).

Another engineering-related issue that emerges in respect of the distinction between extended and embedded cognition centers on the attempt to develop systems with advanced (human-level) intelligence. Suppose that we approach human cognition from an embedded perspective. In this case, the biological brain serves as the point-source for human cognitive capabilities, and it is thus the biological brain that serves as the primary focus of attention when it comes to our attempts to understand the machinery of the mind. The result, from an engineering perspective, is likely to be an approach that focuses on the processing dynamics of the biological brain. Indeed, by building a precise computational replica of the biological brain, we might expect to yield a system with all the cognitive flair, power and sophistication of the human cognitive system. And why not expect this? We have, after all, re-created all the mechanisms deemed to be sufficient for human cognizing.

Now let us look at things from the standpoint of extended cognition. In this case, the mechanisms that are deemed relevant to human cognition are not necessarily confined to the intra-cranial realm. The result is that in attempting to emulate human-level cognitive capabilities, we are obliged to broaden the scope of our computational modeling efforts, taking into account the potentially crucial role played by information processing circuits that extend beyond the neural realm. From an extended perspective, the exclusive focus on neural mechanisms now looks to be a little inadequate. It is almost as if we were attempting to emulate the aeronautical capabilities of a modern fighter aircraft by focusing all our attention on the jet engines. Sure, jet engines are important. And if you build a precise replica of a jet engine, then you will no doubt end up with a system that does at least something interesting. . . perhaps

even valuable. Virtuoso displays of aeronautical expertise are, however, unlikely to be among the list of competences yielded by this particular form of reverse engineering. The same may very well be true of computational modeling efforts that focus all their attention on the biological brain. Crucially, if human intelligence is not realized solely by the biological brain, then a question is raised about the extent to which we can expect large-scale neural simulations [120] and neuromorphic chips [246] to yield information processing capabilities that resemble those of human cognizers.

6

Social Cognition

Thus far, the discussion has centered on how the cognitive profile of individual human agents might be affected by current or future forms of the Web. There is, however, another aspect to this discussion that we haven't touched on as yet. This is the role played by the Web in supporting social interaction and interpersonal communication. The recent growth and expansion of the Web has, to a large extent, been driven by the emergence of systems that alter the shape of our social interactions. Indeed, the advent of social media sites (e.g., YouTube), social networking systems (e.g., Facebook) and microblogging services (e.g., Twitter) has prompted a significant shift in how we think about the Web. Instead of the Web serving as a mere source of information, it is now common to see the Web as something of a 'social space'—a space in which people are able to interact, socialize and share information. We are increasingly witnessing the emergence of what has been dubbed the *Social Web* [see 189]: a suite of applications, services, technologies, formats, protocols and other resources, all united in their attempt to both foster and support social interaction.

The social character of the contemporary Web suggests that it may be relevant to processes that are typically subsumed under the

heading of social cognition [218, 288]. Unlike the kinds of cognition discussed in previous sections (i.e., embodied, extended and embedded), social cognition is concerned with a particular set of processes, namely, those that enable us to make sense of the social world and interact with other agents. In this respect, the Web is an important focus of empirical and theoretical attention. Given the role of the Web in the formation, maintenance and dissolution of inter-personal relationships, it is important that socially-minded philosophers and scientists consider the potential impact of the Web on social cognitive processes.

One widely held view, both inside and outside academia, is that the Web poses a threat to our social cognitive capabilities. One of the main protagonists in this debate is Sherry Turkle. Turkle suggests that we are in danger of substituting the deep and meaningful relationships we have with other human beings for the somewhat more psychologically-insipid forms of engagement we have with digital technologies [515]. A similarly pessimistic view is expressed by Susan Greenfield [206]. She suggests that the overuse of online social networks is contributing to a general decline in our social cognitive skills, perhaps contributing to a rise in the incidence of autistic spectrum disorders. Such claims, it should be clear, are highly controversial, with many arguing that the available evidence is not sufficient to warrant the alarmist character of recent rhetoric [38, 355]. In fact, recent debates about the effects of the Web on human social cognitive capabilities have all the hallmarks of a moral panic. Clowes [111], for example, suggests that the rapid growth and popularity of the Social Web *circa* 2007 outpaced the ability of the scientific community to evaluate its cognitive effects, leading to widespread public anxiety about the Web's impact on socially-relevant cognitive abilities.

In this section, we provide an overview of research relating to social cognition. Although there are many points of interest here, we limit our attention to three topics that have been the focus of recent research. These topics relate to the effect of the Web on person perception and impression formation processes (see Section 6.1), the emergence of so-called social bots within the social media ecosystem (see Section 6.2), and the role of the Web in mediating various forms of social influence (see Section 6.3). Section 6.4 explores an additional topic that is relevant

to social cognition research. This concerns the extent to which the Social Web enables us to exceed the limits on social cognitive processing as imposed by the human brain.

6.1 Person Perception

Person perception (also known as social perception) represents one of the core focus areas for social cognition research. Researchers in this area are concerned with the processes that enable us to form impressions of other people. The cognitive processes of interest are thus ones that relate to the collection and interpretation of information about other individuals [384]. In the standard (face-to-face) case, of course, person perception is informed by a rich variety of linguistic, para-linguistic and non-verbal cues, not all of which are available in the online realm. In spite of this, research suggests that the Web is able to support person perception processes.

One of the main focal points for research in this area relates to the informational cues that are available from social media sites. A notable example is information relating to an individual's social network. As Donath and boyd [149] point out, social networking sites are often set up in such a way that information about a user's social network is automatically broadcast to other users. The number of followers one has on Twitter or the number of connections one has on LinkedIn are two examples of what Donath and boyd refer to as "public displays of connection." Crucially, many of these online social cues are not ones that are (directly) created and maintained by the human users of an online system; instead, the cues are generated by the online system, based on the behavior of the system's user community [508].¹ A Twitter user, for example, has no real control over the information that depicts the number of followers associated with their user account. Instead, this information is computed by the servers that host the Twitter system, and it is the server machines that render this information in graphical or textual form. The process by which online social cues are generated is thus one that relies on patterns of human behavior and the ongoing

¹As a result, such cues are sometimes referred to as "system generated cues" [541].

execution of one or more (machine-based) computational processes. Inasmuch as we see such cues as influencing the processes of person perception, then it should be clear that the outcomes of such processes often have as much to do with the technological properties of the Web as they do with the interpretive capabilities of the Web's human users.

Recent research suggests that online social cues can indeed be used to drive person perception processes, informing social judgements relating to (e.g.) credibility, extraversion, popularity and social attractiveness [508, 528, 541]. Consider, for example, one of the leading social networking sites: Facebook. One of the pieces of information that Facebook provides is the number of 'friends' associated with a user account. This information is generally unavailable (or at least difficult to obtain) in face-to-face contexts, but it is relatively easy to access in the context of the social media ecosystem. Findings by Tong et al. [508] indicate that the availability of this piece of socially-relevant information has a profound effect on perceptions of a profile owner's popularity, extraversion and social attractiveness. By experimentally manipulating the number of apparent friends listed on a Facebook profile page, Tong et al. were able to demonstrate a curvilinear relationship between the number of Facebook friends and perceptions of social attractiveness. In particular, those with moderate numbers of friends were seen as more likeable, relative to those who had a smaller or larger number of friends.

Another study concerning the impact of online social cues on person perception was performed by Westerman et al. [541], this time using the Twitter microblogging system. What Westerman et al. discovered was that the number of 'followers' associated with a Twitter account served to influence judgements relating to the account owner's credibility. As with the study by Tong et al. [508], Westerman et al. report the presence of a curvilinear relationship, with moderate numbers of followers eliciting the highest credibility ratings. Such observations are important because they highlight the ways in which credibility evaluations can be manipulated by relatively simple system-generated cues. Given that the information posted on social media sites may not always be the most reliable, it is important to understand the factors that determine when an individual

is likely to be influenced by such information, based on the perceived credibility of the information source.

These studies are among several examples that illustrate the extent to which online social cues can be used to support person perception processes [519, 527]. Interestingly, the research in this area seems to suggest that person perception occurs quite quickly in the context of social media sites, i.e., people form an impression of another Web user within a relatively short timeframe. This result is somewhat surprising in view of the predictions made by Social Information Processing Theory (SIPT) [526]. One of the central tenets of SIPT is that impression formation and person perception processes occur relatively slowly in an online context, often requiring multiple interactions over an extended time period. Such claims are difficult to reconcile with the results of (e.g.) Westerman et al. [541] who report that users were able to evaluate the credibility of a Twitter user following a single exposure to the user's profile page. "[E]ven in a relatively restricted time period," Westerman et al. conclude, "there are certain types of social information that users can parse and use to make social judgements" (p. 205).

Some insight into the richness of online information, vis-à-vis its ability to support person perception processes, is revealed by studies that attempt to predict personality traits from social media profiles [19, 190, 397]. In one study, Bachrach et al. [19] examined the correlation between personality traits and the properties of a user's Facebook profile (e.g., the size and density of the user's social network, the number of photos the user had uploaded, and the number of groups the user had joined). The results of this study revealed a number of correlations between online social cues and user personality traits. In particular, the best correlations were obtained in respect of extraversion and neuroticism, followed by openness and conscientiousness. By themselves, of course, these results are not sufficient to warrant conclusions regarding the social cognitive utility of online information, since there is no evidence to suggest that human users are able to detect the sort of correlations revealed by statistical methods. Nevertheless, the results do suggest that it is possible, at least in principle, for online social cues to provide the kind of information that could be used to guide person perception

processes. Given that such cues track user behavior over extended time periods, they may very well turn out to be more reliable indicators of (e.g.) user personality than the sorts of cues typically available in face-to-face contexts—and this is despite the fact that face-to-face encounters have traditionally been seen to provide a far richer informational context for person perception.

The process of forming accurate impressions of another individual (the alter) requires the perceiver (the ego) to exhibit a number of cognitive abilities (e.g., the ability to perceive person-relevant information and make inferences based on that information). Interestingly, this emphasis on cognitive abilities provides the basis for an epistemological approach to person perception. In particular, we can view the overall goal of person perception as an attempt to acquire knowledge about (the properties of) other agents. Such knowledge is sometimes referred to as *interpersonal knowledge* or *interpersonal epistemology* [40]. According to Berger et al. [40], interpersonal knowledge is achieved when the ego develops individuating impressions of the alter, including representations of the alter's beliefs. Such impressions are typically formed as the result of an ongoing interaction between the ego and the alter. In particular, Berger et al. cast person perception as an active, temporally-extended process, in which knowledge is acquired as the result of a variety of knowledge-generating strategies, such as interrogation, self-disclosure, deception detection, environmental structuring and deviation testing. Cast in this light, the concept of interpersonal knowledge is one that dovetails with recent work in both cognitive science and contemporary epistemology. From a cognitive science perspective, the emphasis on knowledge-generating strategies establishes a natural point of contact with work that emphasizes the role of actions in helping to support the acquisition of task-relevant information [see 277]. The concept of interpersonal knowledge is also nicely aligned with issues in contemporary epistemology. In particular, the emphasis on cognitive abilities (e.g., the ability to infer person-related properties from a limited set of informational cues) is in perfect accord with virtue-theoretic approaches to knowledge [see 204], especially when it comes to the epistemological position known as virtue reliabilism [203, 392]. A key feature of such

approaches is that knowledge is seen to be grounded in the exercise of an agent's cognitive abilities. In fact, inasmuch as we are prepared to accept that elements of the online environment are actively involved in the realization of person perception processes, then we may have an opportunity to combine a virtue-theoretic approach to knowledge with the notion of extended cognition [see 393]. In this case, we can perhaps see social media systems as working in concert with individual human cognizers to form a hybrid organization whose doxastic outputs reliably track the properties of other Web users.

6.2 Social Bots

The Turing Test [514] has, for many years, been a driving force in AI research. Inspired by the idea of designing a computer algorithm that could pass the Turing Test, researchers have long sought to develop systems with human-like capabilities, especially in the areas of natural language processing and natural language generation. The advent of social media systems has yielded a surprising twist to this decades-long aspirational effort. The twist comes in the form of what are called *social bots* [170]. These are systems that attempt to infiltrate the social media ecosystem, typically by masquerading as human users:

A social bot is a computer algorithm that automatically produces content and interacts with humans on social media, trying to emulate and possibly alter their behavior. [170, p. 96]

In essence, we can think of a social bot as a system that attempts to pass the social media equivalent of the Turing Test. A social bot is thus a system that aims to behave in a manner that is sufficiently human-like so as to be indistinguishable from a genuine human user. This is obviously a different sort of challenge to that posed by the original Turing Test. In particular, the nature of the formats and protocols adopted by social media systems may make it easier for a social bot to emulate the behavior of human users (consider, for example, how the 140 character limit in Twitter may simplify the task of generating suitably convincing linguistic content). In addition, the fact that social

bots are embedded within the social media environment provides an interesting set of opportunities for behavioral emulation. Social bots may, for example, monitor the behavior of Web users and attempt to model their behavior based on the insights they glean from the resultant data sets. They may also access the online profiles of Web users and attempt to imitate or impersonate those users [187]. Finally, social bots may attempt to mask their true identity by exploiting Web-accessible content. As noted by Ferrara et al. [170]:

...social bots can search the Web for information and media to fill their profiles, and post collected material at predetermined times, emulating the human temporal signature of content production and consumption—including circadian patterns of daily activity and temporal spikes of information generation. (p. 99)

One reason why social bots are interesting in the context of social cognition research is because issues of social sensemaking, impression formation and person perception lie at the heart of the original Turing Test. In order to pass the Turing Test, a synthetic agent needs to behave in a manner that is suitably aligned with the social cognitive capabilities of a typical human interlocutor. Inasmuch as the social bot achieves this goal, then it is likely to be regarded as a human agent. This is important, because in an online context such perceptions are apt to alter the cognitive and behavioral responses of large swathes of the human user community. Indeed, much of the recent research on social bots has focused on the potential threat they pose to the online social environment.² In the aftermath of the 2013 Boston bombing incident,

²It should also be noted that social bots pose a threat to the integrity of empirical research that seeks to use the Web as a platform for social and cognitive research. The failure to disambiguate social bots from genuine human users may, for example, undermine the validity of research that seeks to understand the dynamics of collective cognition (see Section 7). A related point concerns the attempt to monitor and predict worldly events using social media systems. The real-time analysis of Twitter tweet streams, for example, has been used to predict or detect the outbreak of epidemics [296], election results [513], and earthquakes [419]. The extent to which these predictive efforts are compromised by the presence of social bots is, at the present time, unclear.

for example, a number of false accusations were widely circulated on Twitter, mostly as the result of social bots automatically retweeting existing posts [213]. The widespread dissemination of this sort of false information is apt to nullify the otherwise positive role that social media systems may play with respect to the performance of collective cognitive tasks (e.g., collective sensemaking tasks) [284, 363] (see Section 7.1). A related concern arises in respect of political campaigns. As noted by Ratkiewicz et al. [402], social bots are a potential threat to democracy, in the sense that they may artificially inflate support for a political candidate and thus influence the outcome of political elections. What is interesting here, is that social bots are perceived to be a threat precisely because they are apt to be confused with human users. This suggests an interesting reversal of the challenge posed by the original Turing Test. The current difficulty, it seems, is not so much the development of synthetic agents that are able to emulate human behavior; rather, it is the development of algorithms that are sufficiently discerning as to discriminate *bona fide* human agents from their more ersatz counterparts [170]. The upshot of all this is likely to be an ‘arms race’, in which advances in bot detection methods are paralleled by the attempt to develop ever-more sophisticated (i.e., human-like) social bots. Predicting the outcome of this arms race is difficult; however, one possibility is that (in the specific socio-ecological niche of the Web) human and synthetic agents will be, in large part, indistinguishable. Future social bots may thus be treated as kindred social cognizers of the online social realm.

6.3 Social Influence

Among the many topics that are of interest to social psychology, social influence is perhaps one of the most important. As noted by Mason et al. [329], “The study of social influence—the ways other people affect one’s beliefs, feelings, and behavior—in large measure defines social psychology” (p. 279). Much the same, of course, could be said about disciplines whose focus of interest lies above the level of individual agents, e.g., the disciplines of sociology, economics and political science.

Social relationships are obviously key to understanding the nature of social influence; for they are the conduits through which social influence is often exerted. Every relationship that we establish with others is thus a channel that supports the flow of information, enabling our thoughts and actions to influence the thoughts and actions of others, and enabling our own thoughts and actions to be influenced in return. It is partly for this reason that the study of social influence often goes hand in glove with the study of social networks. By understanding something about the structural organization of social networks, scientists hope to acquire insight into how the properties of both individuals and groups are affected by issues of social structure [61, 95, 446, 555]. Such goals are not, of course, exclusive to social science, and the scientific study of networks—the basis of the nascent discipline of network science [25, 121]—is one that is common to a number of scientific disciplines [24, 68], including the disciplines of both cognitive science [26] and Web Science [504].

The Social Web is perhaps the most obvious point of interest for those concerned with issues of social influence. Social Networking Sites (SNSs), such as Facebook, for example, provide a relatively new means by which social networks can be managed, maintained and sometimes modified, and we would therefore expect such systems to play an important role in social influence. Of particular significance is the fact that many SNSs provide APIs to support the derivation of population-scale data sets, yielding insight into the time-variant structural profile of online social networks and supporting studies into the factors that regulate social influence.³ One example of a study that combines the use of network analytic techniques with an attempt to monitor social influence in an online SNS is provided by Aral and Walker [13]. Aral and Walker monitored the tendency of Facebook users to adopt a commercial Facebook application in response to messages sent from a peer in their local social network. The subsequent analysis of data from 1.3 million Facebook users yielded a number of interesting results, including the finding that younger users are more susceptible to influence than

³As is evidenced by the comments of those who work in computational social science [123, 185, 302, 485], such data sets are often seen to provide an opportunity to improve our understanding of social processes.

older users, men are more influential than women, women influence men more than they influence other women, and married individuals are the least susceptible to social influence. These results, it should be clear, reveal something about the moderating effect of individual, agent-level properties—e.g., age, gender and marital status—on the influential status of a given individual. But in addition to this, Aral and Walker suggest that the structural organization of a social network is also relevant to social influence. In particular, the connections between influential individuals appears to accentuate their influential capacity. This suggests that when it comes to product adoption decisions on Facebook, the tendency of a group of influential friends to act in concert with one another may be a significant factor in determining the extent of social influence [13].

It is natural to think of social influence as something that occurs in a ‘peer-to-peer’ fashion, with one agent influencing the thoughts and actions of their network neighbors (e.g., friends on Facebook). But these are not the only forms of social influence that can be found on the Web. Web-based systems thus provide a rich array of (socially-derived) informational cues that can serve as the basis for indirect forms of social influence. To help us see this, consider the kinds of cues that were discussed in Section 6.1. These cues (e.g., the number of Twitter followers associated with a user account) represent information about a *social* property, and we have already seen that such cues are apt to influence the thoughts and actions of the average Web user. What we seem to have here is a form of influence that connects the realm of individual thought and action to the realm of social or collective phenomena (i.e., phenomena that apply to collections of individuals). As another example of this sort of influence, consider the way in which YouTube provides an indication of the popularity of uploaded digital resources (i.e., videos) by explicitly representing information about the number of times such resources have been viewed (or downloaded) by other users. This particular informational cue provides an explicit representation of an aggregate measure of collective behavior (i.e., it is a cue that provides information about a collective phenomenon). In addition to this, however, it is also a cue that human Web users can perceive and

process. In other words, the cues in question are ones that are apt to mediate the cognitive and behavioral responses that individuals have in response to specific aspects of the online social environment.

One way of understanding these sorts of social influence is via the sociological notion of *immersion* [12, 122]. This is a particular form of social influence that occurs in response to information about social or collective phenomena. In a Web-based context, immersion occurs when individual behavior is influenced by the presence of representations that refer to patterns of collective behavior, or (more generally) some property that is pitched at the social level of analysis.⁴

- Access to information about the social position of others (e.g., the number of followers on Twitter or the number of friends on Facebook) can impact person perception processes [508, 541]. Such processes are poised to affect the responses directed towards other individuals (see Section 6.1).
- Cues indicating the popularity of certain kinds of ‘deviant’ behavior can help to normalize those behaviors and facilitate their more general expression. Examples, in this case, include the normalization of self-harm behavior through social media [162] and the use of social frequency information as a means of justifying cybercriminal activity [486]—the so-called “everyone else is doing it” rationalization [358].

⁴This is similar to the forms of influence proposed by social scientists in attempting to understand the relationship between social phenomena and individual action. Social phenomena (which arise as the result of social interaction) are thus seen to shape and constrain the actions of the individual agents that are (in some way) exposed to the phenomena. This sort of idea is commonly encountered in the social science literature. In their discussion of causal mechanisms in the social sciences, for example, Hedström and Ylikoski [227] suggest that progress in our understanding of social phenomena requires the identification of “situational mechanisms by which social structures constrain individuals’ action and cultural environments shape their beliefs and desires” (p. 59). In essence, we can see the Web as providing the means by which social phenomena are able to effect changes in individual thought and action. The following are some of the more obvious examples of immersion in an online context:

- The availability of cues concerning the social popularity of certain actions, topics, content and so on can, in some cases, provide the sort of influence that is unable to be exerted through more direct, inter-individual forms of communication. Trending topics on Twitter are one example of this sort of influence [see 199].
- Traces of past behavior (e.g., purchasing decisions) can sometimes influence future behavior. An example is provided by collective attention effects in YouTube [250]. In this case, the future productivity of individual users is influenced by explicit representations of the number of times a user's videos are downloaded.

From the perspective of social cognition, cues that provide information about collective cognitive (or socio-cognitive) phenomena are of particular interest. Such phenomena concern the cognitive responses of *multiple* individuals as expressed through (e.g.) aggregate measures. Consider, for example, cues that track levels of collective interest or collective attention in an online resource [32]. Prominent examples are the number of times a particular resource has been read (as in the case of academic paper citations) or the number of times a particular resource has been downloaded (as in the case of YouTube videos). As we have seen, the information provided by such cues can exert a significant influence on one's future patterns of behavior [see 250]. All this, we suggest, leads to a specific hypothesis regarding the cognitive impact of socio-cognitive cues. For the sake of convenience, we will refer to this as the *socio-cognitive feedback hypothesis*:

Socio-Cognitive Feedback Hypothesis

Informational cues that track the cognitive responses of other users (e.g., indicators of collective attention) play an important role in determining the nature of one's future interactions and engagements with the online realm. This is especially so when it comes to the social evaluation of one's own contributions.

There is, of course, nothing that is particularly radical about this hypothesis: it is simply a reminder that the Web can, on occasion,

help to shape human behavior by virtue of its ability to confer social rewards and mete out social punishments. The socio-cognitive feedback hypothesis thus serves as an important reminder of the power and potential of the Web as a source of information to which we, as a social species, are exquisitely (and sometimes painfully) attuned. It is well known, for example, that social recognition and social status are valuable commodities that can be used to incentivize human behavior in lieu of financial remuneration [249], and this perhaps helps to reveal one of the factors that motivates participation in online systems [295]. Indeed, it is the perhaps the capacity of the Web to deliver social rewards that explains why we humans are so inclined to devote ever-increasing amounts of time and energy to it.

The socio-cognitive feedback hypothesis is of particular interest and relevance in the context of the Web. One of the reasons for this relates to the fact that the kind of information that can be acquired in the online realm is not the same as that provided by conventional forms of face-to-face social interaction. As was mentioned above (see Section 6.1), the Web provides a means of tracking and explicitly representing socially-relevant information (e.g., details of one's social network), and this information may not be so readily available in an offline (i.e., face-to-face) context. In addition, it should be clear that the kind of informational cues we are dealing with here are relatively easy to manipulate (imagine a state-of-affairs in which the number of downloads for a particular YouTube video was artificially elevated by a warm-hearted hacker). Setting aside the ethical quandaries associated with such forms of deliberate manipulation (of which there are no doubt many—both pro and con), it follows from the socio-cognitive feedback hypothesis that the manipulation of socially-relevant information (e.g., online social cues) may, at least in some cases, yield patterns of user behavior that are commensurate with the goals of an online system (or at least the goals of the system's corporate sponsors).

6.4 Social Brains, Extended Minds

Recent theories concerning the evolution of the human brain have focused their attention on the challenges and constraints imposed by the

human social environment. Social cognitive capacities often lie at the heart of such theories. The general idea is that in the attempt to deal with the vagaries of human social life, our species evolved specialized skills of social cognition [237]. The Machiavellian intelligence hypothesis [72, 546], the social intelligence hypothesis [287], and the social brain hypothesis [155, 156, 158], are some examples of theoretical proposals that appeal to this sort of idea, or at least some variant thereof.

One of these proposals, we suggest, is of particular interest and relevance to Web Science, especially when it comes to a consideration of the Social Web. This is the social brain hypothesis. The social brain hypothesis assumes that some aspect of brain size determines, or sets limits on, core social cognitive abilities. These limits, in turn, determine the number of relationships (and hence group size) that social animals can effectively maintain as coherent entities. Such constraints yield a species-specific cognitive group size [158], which in the case of our own species is postulated to be around 150 [157]. This is the so-called *Dunbar's Number*.

Relative to the constraints on social cognitive capacity imposed by the social brain hypothesis, the advent of the Web (and especially the Social Web) raises an important question: Is it possible for the Social Web to support a form of cognitive enhancement that surpasses the neurologically-determined performance limits associated with the 'social brain'? This question strikes a chord with the earlier discussion on extended cognition (see Section 4). For if we accept that extended cognition provides a means by which human agents are able to expand the cognitive reach of the biological brain, then perhaps the Social Web can be seen to yield an extended cognitive capacity that transcends the neurological limits imposed on social cognitive processing. In other words, perhaps by factoring in the resources of the Social Web, we humans are able to deal with social environments of far greater complexity than would otherwise be suggested by the social brain hypothesis.

For the sake of convenience, let us refer to this idea regarding materially-extended social cognitive capacities as the *extended social intelligence hypothesis*:

Extended Social Intelligence Hypothesis

The advent of the Social Web transforms the social cognitive capabilities of the biological brain. By virtue of our ability to incorporate the resources of the Social Web into our cognitive processing routines, we are able to enhance our capacity to process social information. Relative to the Social Web, the capacity for cognitive extension leads to an inflation of social cognitive capacities: the *Social Web-extended mind* is (perhaps) a supersized social cognizer.

The extended social intelligence hypothesis proposes that some of the cognitive mechanisms that underpin our processing of social information are realized by circuits that extend beyond the neural realm. Inasmuch as this is the case, then there seems little reason to assume that the social cognitive limits of the human individual are always the same as those mandated by the social brain hypothesis.

A number of commentators have voiced opinions consistent with this idea. Wellman [540], for example, suggests that “social media have increased the carrying capacity of relationships, with heavy internet users having more close ties” (p. 174). Similarly, de Ruiter et al. [416] note the potential of contemporary technologies to expand our social horizons, extending the reach of our social cognitive capabilities beyond their neurologically-determined limits:

Dunbar’s assumption that the evolution of human brain physiology corresponds with a limit in our capacity to maintain relationships ignores the cultural mechanisms, practices, and social structures that humans develop to counter potential deficiencies. [416, p. 559]

de Ruiter et al. go on to suggest that:

The human neocortex may be finite, but human capacity appears malleable and expansive in different sociocultural and technological contexts. [416, p. 562]

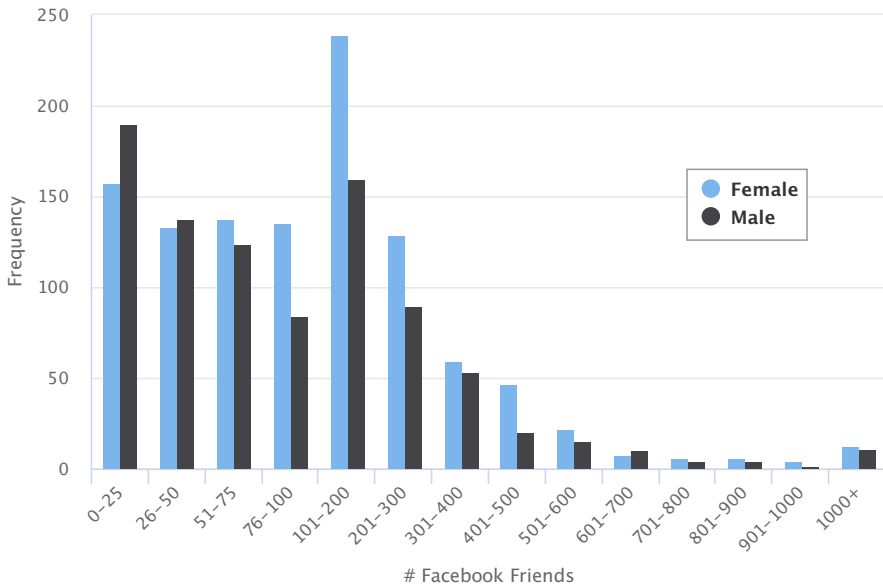


Figure 6.1: Frequency distribution of reported number of Facebook ‘friends’ for male and female respondents ($N = 2000$) (source: Dunbar [160]).

These claims regarding the augmentative potential of the Social Web have been disputed by Dunbar [159, 160]. In one study, Dunbar [160] examined the number of friends listed on Facebook as a means of evaluating the claim that social media allow us to increase the size of our social networks. In fact, he found that the size of online social networks resembled those of offline social networks (see Figure 6.1). On the basis of this result, Dunbar concludes that cognitive constraints pertaining to social network size are not alleviated by the adoption of social media technologies:

... on the evidence available so far, it seems unlikely that the digital media will significantly change our social lives, at least in terms of the number and intensity of different kinds of relationships. Our cognition, inherited as it is from our primate ancestry, seems to make that impossible. [159, p. 2199]

Despite the negative result, the issue of whether the Social Web can enhance aspects of social cognitive functioning remains an important

one for the philosophical, scientific and engineering community. In particular, future work in this area should aim to determine whether Web technologies are able to improve the processing of social information as the result of cognitively-potent forms of bio-technological merger. One reason for continued research in this area relates to the changing nature of our contemporary society, as captured by the notions of *network individualization* [400] and the *network society* [85, 144, 400]. According to the proponents of these concepts, we are moving away from an era in which society is organized around fixed and relatively enduring collectives (such as families, traditional friendship circles, organizations and communities) and entering an era in which society is increasingly organized along the lines of multiple, overlapping social networks. One consequence of this transition is that individuals are under increasing pressure to form more fluid and flexible social networks as a means of satisfying basic needs, such as those pertaining to the social, emotional, sexual, physical, and financial aspects of human life.

Inasmuch as we accept all this, then it seems that social media and social networking systems may be able to play an important role in helping our species adapt to our new socio-ecological niche. Just as the aqueous environment of dolphin society is believed to lie at the heart of their tendency to form highly dynamic and fluid task-oriented social groups [401], so perhaps we can think of the Social Web as an environment that supports our ability to form fluid, flexible and multiply overlapping social networks. Perhaps, in other words, the Social Web is just that part of the online cognitive ecology [460] that enables us to cope with the ever-shifting sands of the human social world.

7

Collective Cognition

Ever since the advent of Web 2.0, which is characterized by greater levels of user participation in the creation, maintenance and editing of online content, the Web has provided ample opportunities to support socially-distributed information processing. This highlights the potential role of the Web in supporting episodes of collective and specifically distributed cognition, i.e., cognition in which the relevant cognitive processes (e.g., reasoning, remembering and problem-solving) are distributed across a mixture of multiple individuals and technological artefacts (see Section 2.3.2). A considerable amount of interest has been expressed by the Web Science community in this particular form of cognition [90, 91, 92]. Previous work, for example, has studied the Web in relation to a variety of socio-cognitive phenomena (the collective variants of individual-level cognitive phenomena). This includes work on collective problem-solving [350], collective sensemaking [284], collective attention [32], collective creativity [557] and collective memory [367]. Issues of distributed cognition also lie at the heart of a number of prominent research areas relevant to Web Science. This includes work relating to social computation [270], human computation [300, 348, 350],

398], collective intelligence [59, 215, 322, 323], social machines [236, 382, 466], and the global brain [47, 238].

Given the quantity of Web Science work that has been undertaken in relation to collective cognition, it is unlikely that any single review will be able to do justice to the existing literature in this area. The current review is, of course, no exception. Rather than attempt to provide a comprehensive overview of existing work, the current section focuses on a number of topics that are likely to be of common interest to those working in Web Science, cognitive science and the philosophy of mind. Such topics include the relationship between network structure and collective cognitive performance (see Section 7.1), the relationship between social feedback and certain forms of collective intelligence (see Section 7.2), the role of technology-mediated social participation in the construction of cognitively-potent informational ecologies (see Section 7.3), the status of some Web-based socio-technical systems as problem-solving organizations (see Section 7.4), the status of Web-based systems as the engines of knowledge discovery (see Section 7.5), and the extent to which the social properties of the Web support socially-extended forms of cognition (see Section 7.6). The section concludes with a discussion of mandevillian intelligence (see Section 7.7). This is a specific form of collective intelligence, in which cognitive constraints, limitations and biases at the level of individual agents are seen to play a positive functional role in yielding collective forms of cognitive success.

7.1 Socio-Cognitive Circuits

In order to support effective forms of collective cognition on the Web, it is important to develop a better understanding of the forces and factors that shape collective cognitive outcomes. One factor that has been the focus of considerable research attention is the structure of the communication network in which agents are embedded [270, 328, 330, 331, 461]. Network structure is important because it influences the dynamics of information processing. Communication networks are thus somewhat similar to neural networks or electrical circuits, in that the topological structure of the network—the pattern of connections between

nodal elements—can be seen to exert a significant influence on the kind of information processing implemented by the networked ensemble.

It has long been known that network structure is important to at least some forms of collective cognizing. Classic studies by Bavelas [33] and Leavitt [303], for instance, revealed that the structure of a communication network affects the performance profile of problem-solving groups, with more centralized structures yielding the best results. Subsequent research yielded a twist to these early findings, showing that the relation between network structure and task performance is modified by task properties [436]. It is thus not the case that a particular kind of network structure is best for *all* kinds of tasks (or variations of the same task): different networks exert different effects depending on the nature of the task that is being performed.

One task property that seems to be relevant to collective performance is task difficulty, i.e., the difficulty of the problem that must be tackled by a community of problem-solving agents [330, 436]. When subjects are confronted with a simple problem, it seems that the best network structure is one that connects every individual to every other individual, thereby yielding a fully-connected network. On more complex problems, however, it seems that more limited forms of connectivity are preferable. The reason for this is not always clear, although limited connectivity may help to impede the rate of information flow between individuals and thus prevent premature convergence on sub-optimal or inaccurate solutions [253, 301, 455, 461].

Studies using social computation experiments have also contributed to our understanding of the relationship between task properties and network structure [270]. One of the findings to come out of this body of research is that even subtle differences between tasks can influence the performance outcomes associated with different network structures. In one study, for example, Judd et al. [264] compared the impact of multiple network structures across two tasks. They found that network structure elicited opposing behavioral effects in the two tasks, and this was despite the fact that the tasks were judged to be somewhat similar from a cognitive perspective. These results lend support to what we will call the *task specificity hypothesis*:

Task Specificity Hypothesis

Different network structures (e.g., small-world, fully-connected, etc.) yield differential effects on collective cognitive outcomes based on the properties of the cognitive task that is being performed.

For the most part, studies exploring the effect of network structure on collective cognition have been undertaken in controlled laboratory environments where it is relatively easy to impose limits on inter-agent communication. When it comes to the Web, however, it is much more difficult to determine the structural profile of the socio-cognitive circuits associated with specific episodes of collective cognitive processing. In fact, it might be thought that the Web acts to remove barriers to information flow, enabling agents to share information in an environment that is (more-or-less) accessible to all agents. Once an item of task-relevant information has been posted online, for example, it becomes available to everyone who accesses the online environment, and this suggests that the Web might act to reduce the path length of a socio-cognitive network. In other words, the Web seems to afford the possibility for highly efficient and widespread forms of information dissemination—the forms of information dissemination that are typically associated with fully-connected networks (see Figure 7.1).

This vision of the Web as yielding something akin to a fully-connected network presents us with a potential dilemma. For inasmuch as we embrace the task specificity hypothesis, then it should be clear that the Web is likely to be suited to some kinds of tasks (i.e., those benefiting from fully-connected network structures) but not others. In fact, the extent to which the Web can exert a positive effect on collective cognition is still far from clear. One problem is that rapid and widespread forms of information dissemination may lead to situations in which the creativity of a group is undermined. Consider, for example, the phenomenon of production blocking, which occurs in group solving situations, specifically brainstorming sessions [142]. Production blocking is the tendency for the contributions of one individual to block or inhibit contributions from other group members, resulting in reduced levels of creativity relative to what might have been expected if the group

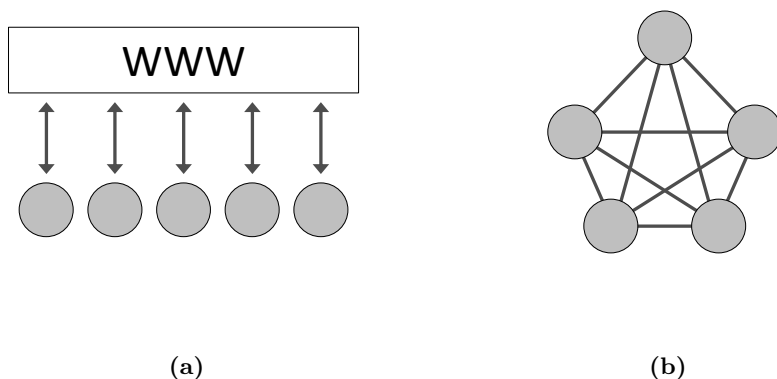


Figure 7.1: The status of the Web as a shared information space (a) yields the same pattern of information flow as that encountered in the case of a fully-connected network (b).

members had worked independently. It thus seems that the form of information dissemination enabled by the Web might act to undermine the *collective* creative potential of a group of agents: instead of stimulating a greater number and diversity of ideas, the Web might work to impede rather than improve the creative process.

Much the same could be said of another form of collective cognition, namely, collective sensemaking.¹ In this case, rapid forms of information dissemination have been shown to impair performance, especially in situations where the information to be processed is conflicting, ambiguous or incomplete [253]. This may help us understand why Web-based forms of collective sensemaking sometimes go awry. Consider, for example, the use of Reddit to support large-scale collective sensemaking efforts in the case of the 2013 Boston bombings [363]. Despite providing a potentially

¹Collective (or team) sensemaking is a specific form of collective cognitive activity, in which multiple individuals attempt to resolve the features of objects or situations based on incomplete, ambiguous and/or conflicting information [278, 467]. Note that the meaning of the term “sensemaking” in the present context is somewhat different to that typically encountered in the philosophy of mind and cognitive science literature [e.g., 502].

useful platform for users to upload and analyze images, identify potential suspects, and highlight topics for further analysis, the user community ultimately failed to identify the bombing suspects.² This failure is disappointing in view of the wealth of material available, and it highlights the need for further work to improve our understanding of the factors that affect the collective cognitive performance of online communities.

It is in the attempt to improve our understanding of collective cognition that we encounter an important point of contact between Web Science and cognitive science. In particular, the Web provides us with a valuable opportunity to observe the dynamics of information flow within a collective cognitive organization. In Section 4.7, we encountered the idea that the Web enables us to monitor an individual's cognitive processing routines—to “see the mind in action,” as Krueger [286] puts it. A similar idea can now be rolled out in respect of collective cognition. The basic idea is to use the Web as an observational platform to track the cognitive consequences of particular forms of (online) information flow. The goal of such observational efforts resembles those of digital cognitive ethnography [257], which is a well-established technique to study the behavior of small-scale (typically team-based) socio-technical systems. One example of digital cognitive ethnography is provided by Hutchins et al. [258]. In this case, digital cognitive ethnography was used to study the socio-technical interactions involved in piloting the Boeing 787 Dreamliner aircraft. Hutchins et al. describe how the use of digital recording technology yielded a rich array of performance-relevant behavioral data, including “digital audio recordings, digital pen data (recording notes made by the pilots as well as notes made by observers), wearable eye tracking on both pilots, and digital data from the simulator itself” [257, p. 43].

Of course, the data obtained from the Web is unlikely to be as rich or as detailed as that provided by digital cognitive ethnography. Nevertheless, there is evidence of such methods being used to support the analysis of Web-based systems. In their analysis of distributed cognitive processes in Wikipedia, for example, Geiger and Ribes [184] discuss the use of *trace ethnography*, which they describe as “a novel method for

²See <http://www.bbc.co.uk/news/technology-22214511>.

studying the complex interactions that occur in sociotechnical systems.” The further development and application of such methods may be useful in helping us gain a better understanding of the socio-cognitive machinery associated with certain kinds of Web-based system.

7.2 Social Feedback

In addition to issues of network structure, the access that individuals have to information about the progress or status of some form of collective cognitive processing has also been the focus of research attention. This is of particular interest in a Web Science context, because the Web provides a unique opportunity to gather and exploit information about the judgements, decisions or activities of other individuals (see the discussion of the socio-cognitive feedback hypothesis in Section 6.3). Consider, for example, a situation in which a group of Web users are asked to derive an estimate of some unknown parameter. For the sake of argument, let us assume that users are asked to estimate the number of crimes that are recorded annually in the city of London. In this situation, the statistical average of everyone’s estimates should approximate the actual number of crimes recorded, and this exemplifies one way in which a system like the Web may be used to exploit what is commonly referred to as the *Wisdom of Crowds* [490]. Imagine, however, that a Web-based system that mediates this particular instance of collective intelligence provides feedback on the estimates made by other users. How will the provision of this information influence user behaviour? And, in the long-term, will the feedback lead to better or worse performance relative to what might have been expected in situations involving the absence of feedback?

In order to answer this question, Lorenz et al. [313] devised an experiment in which participants were asked a number of questions—the answers to which were not known in advance by any one individual. They then manipulated the level of feedback that participants were given about the responses of other participants across a number of trials. The results revealed that feedback often works to undermine collective performance. Rather than being able to derive estimates that were, at the collective level, close to the actual answer, subjects given high

levels of feedback settled on responses that were, again at the collective level, worse than those seen in situations where subjects were given no feedback at all. In accounting for these results, Lorenz et al. [313] posit a *social influence effect*, in which the feedback about subject ratings is deemed to progressively reduce the diversity of ratings within the group without a corresponding improvement in group-level accuracy. These results suggest that although the Web provides an environment in which a variety of kinds of information can be gathered during the course of socially-distributed information processing, not all of this information *should* be made available to the individual agents engaged in the process. Instead, the results call for a more nuanced approach: one in which the system works to adaptively regulate the availability of different kinds of information in ways that are sensitive to the nature of the task being performed. In essence, what is required here is a way of dynamically organizing the setup of Web-based socio-technical systems in order to meliorate group-level cognitive processes within a variety of different task contexts. As our understanding of the factors that govern collective cognitive success improves, so it is likely that systems will emerge in which the pattern of information flow and influence between problem-solving agents is deliberately monitored and controlled so as to improve the prospects for collective cognitive success.

7.3 Collaborative Construction

In some cases, the processes by which specific informational ecologies come into being on the Web are glossed in distributed or collective cognitive terms. Consider, for example, the way in which the linking behavior of Web users yields a body of information that can be used to support the operation of Web search engines. The most popular example is perhaps the PageRank algorithm, as used by Google Search [66]. Here, the editing actions of countless numbers of Web users serves as the analytic substrate for machine-based processes that seek to enhance the accessibility of online information. Heintz [233] suggests that we should see such processes as a specific form of distributed cognition:

An essential aspect of second generation search engines is that their ranking algorithms take as input the linking behaviour of web-users. The consequence is that search engines together with web-users constitute a distributed cognitive system for the attribution of reputation, visibility, and, eventually, credibility. [233, p. 388]

Appeals to the notion of distributed cognition have also been made in respect of collaborative tagging [478]. Collaborative tagging refers to the process by which descriptive ‘tags’ are added to online resources as a means of supporting subsequent search and retrieval efforts. Prominent examples include sites such as de.licio.us³ (for websites) and Flickr⁴ (for photographs). Such forms of collaborative editing and enrichment are of particular importance when it comes to non-textual media resources, such as images, videos and audio resources. Given the inherent difficulties in using automated techniques to resolve the semantic content of such resources, the ability to rely on the human user community is of obvious benefit in addressing what is sometimes referred to as the “semantic gap,” a term that refers to “the discrepancy between low-level technical features of multimedia, which can be automatically processed to a great extent, and the high-level, meaning-bearing features a user is typically interested in and refers to when searching for content” [442, p. 37].

Collaborative tagging is perhaps best thought of as a form of enrichment activity—it is a way of enhancing the structure of an existing informational ecology in a manner that (hopefully) supports all manner of subsequent cognitive activities. But there is another sense in which issues of collective cognition have been deemed relevant to the process of collaborative construction. This emerges in respect of recent attempts to understand the *mechanisms* that underlie the generation of Wikipedia content. In particular, the creation and editing of Wikipedia articles has been characterized as a form of human stigmergy [239, 385], similar to the forms of stigmergy exhibited by eusocial insects.⁵ To help

³<https://del.icio.us/>

⁴<https://www.flickr.com/>

⁵The notion of stigmergy is typically encountered in the fields of swarm intelligence [60] and swarm cognition [509, 516]. The term was originally used to account

us understand this claim, it will be useful to consider how stigmergic processes operate in the case of swarm cognition [509, 516]. Swarm cognition is a term that is applied to species of eusocial insects, such as termites, ants, bees and wasps. It refers to the surprising intelligence that is exhibited by these species in solving certain kinds of problem. The ability of termites to construct architecturally elaborate termite mounds, for example, has been linked to the operation of stigmergic processes: as one termite drops a mud-ball, so it leaves a pheromone marker that encourages other termites to deposit mud-balls nearby, and as the collection of mud-balls increases in size, so macro-level architectural structures begin to emerge as the result of collective, pheromonally-mediated behaviors [516]. A similar sort of explanatory account emerges in respect of the foraging behavior of ants. In this case, the ability of ants to coordinate their collective efforts is deemed to rely on the deposition of behavior-modifying scent trails [489]. What is common to these (and other) cases of swarm cognition (or swarm intelligence) is the idea that the solution to some (otherwise intractable) problem emerges as the result of individual responses to a collectively-configured environment.

The parallels with Wikipedia now start to become clear. For one way to think about Wikipedia is as a collectively-configured environment in which individual editing actions occur in response to the (sometimes subtle) changes made by other users. As each user edits a Wikipedia article, so they alter the online environment in ways that shape, stimulate and support the editorial responses of other users. Played out over sufficient time, such processes lead to the informational equivalent of the termite's mound: a complex informational edifice that no individual, working by themselves, could ever hope to construct.

It is thus by virtue of the appeal to stigmergic mechanisms that we are able to identify an important link between two ostensibly disparate areas of research: the study of eusocial insects on the one hand,

for the capacity of certain species of eusocial insects to create complex structures (e.g., termite mounds) and exhibit complex behavioral patterns (e.g., trail following). A useful definition of stigmergy is provided by Heylighen [239]. He suggests that stigmergy is “an indirect, mediated mechanism of coordination between actions, in which the trace of an action left on a medium stimulates the performance of a subsequent action” (p. 6).

and the collaborative construction of online, Web-based resources (e.g., Wikipedia) on the other. One reason this link is important is because it helps to reveal commonalities in the mechanistic underpinnings of two forms of collective cognition, namely, swarm cognition and distributed cognition. Another reason the link is important is because it helps to highlight the role of the environment in yielding certain forms of collective cognitive success. In the case of the termite mound, for example, the interaction between structural, atmospheric and hydrological parameters plays a crucial role in fixing the functional significance of individual responses [516]. Similarly, when it comes to Wikipedia, we can see the properties of the online environment—the way in which Wikipedia presents information to its user community—as playing a crucial role in enabling the larger socio-technical organization (Wikipedia+users) to produce a successful outcome. What is crucial, here, is the idea of some form of adaptive alignment between the properties of the online technical environment and the cognitive/behavioral dispositions of the human user community. When such forms of alignment are established, the online technical system can be seen to work in concert with the human community, yielding (in the case of Wikipedia) a rich and reliable source of information that is relevant to a broad array of epistemic activities [166, 167].

All this talk of collectively configured environments and the ostensible link between distributed and swarm cognition helps to reveal a couple of additional points that are likely to be of common interest to Web Science, cognitive science and the philosophy of mind. The first point relates to issues of social scale. In particular, large-scale forms of social participation involving thousands (if not millions) of individuals are a common feature of many forms of Web-based distributed cognition. In the case of Wikipedia, for example, tens of thousands of individuals participate in the editing of online articles, and Google Search requires even larger numbers of individuals (hundreds of thousands to millions). Crucially, large-scale forms of social participation are often critical to the success of these systems. In the absence of large-scale social participation, for example, Wikipedia could not have the coverage it does, nor could it update its articles in a timely fashion. (Both of these features, it

should be clear, are important when it comes to Wikipedia's status as a source of epistemically-significant information [167].) It is here that we encounter a potentially important difference with the kinds of systems that are the traditional empirical targets of distributed cognitive science; for such systems—the classic case being the U.S. Navy ship studied by Edwin Hutchins [254]—seldom exceed more than a few hundred individuals. In this sense, systems such as Wikipedia seem to have more in common with the entomological systems studied by the proponents of swarm cognition and swarm intelligence. Such forms of correspondence are interesting when one considers the supposed scale-dependent nature of different social coordination protocols (i.e., the possibility that specific social coordination protocols require a minimum group size in order to be effective) [e.g., 389].

Another point to note is that with all the cases of collaborative construction reviewed above, human collective efforts are helping to shape the structure of the online informational environment. There is, of course, nothing that is particularly remarkable about this claim: every form of interactive engagement with the Web (especially since the advent of Web 2.0) is likely to alter the online environment in some way or another, even if the changes simply relate to the imprints of our digital footprints as we traverse the online realm. But beneath the seeming banality of this claim, there lies something of crucial significance. And once again, it is the domain of entomology that serves as the source of inspiration. The point of interest comes in respect of the way in which processes of collaborative construction in eusocial insects have been glossed as a form of *ecosystem engineering* [see 136]. This establishes a natural point of contact with recent work in the philosophy of mind concerning the notions of *cognitive niche construction* [102] and *ecological engineering* [480]. The basic idea is that, inasmuch as we regard the Web as a resource that influences the course of our individual and collective cognitive endeavors, then perhaps our engagements with the online environment can be viewed in similar terms, i.e., as a form of cognitively-relevant ecological engineering [see 460]. Even if we demur from the conclusion that the processes by which online informational ecologies come into being should be counted as *bona fide* cases of

collective or distributed cognition, we can surely accept the idea that many of our forays into the online realm lead to the creation and configuration of an environment that is poised to exert an ever-more pervasive influence on the things we think, feel and do.

7.4 Social Machines

Many of the most popular systems on the Web today are ones whose content is provided by their user communities. In the case of Wikipedia, for example, the majority of the online content is not provided by a select group of system developers—as might very well have been the case prior to the advent of Web 2.0. Instead, the majority of the content is supplied by the actual users (or, more accurately, the editors) of the Wikipedia system. System developers, of course, still have an important role to play here; for it is they who provide the infrastructure that enables the user community to supply the actual content. Nevertheless, it is clear that something has changed since the early days of the Web. Instead of pre-populating an online system with content in advance of its actual use, we have witnessed the rise of an alternative development model: one where the process of content generation is delegated (or outsourced) to the user community. As is evidenced by the likes of Wikipedia, Twitter, Facebook, YouTube and Flickr, this approach has turned out to be extraordinarily successful. In the absence of user contributions, it might be difficult (not to mention costly) for online systems to garner the sort of content that maintains their popularity (and sustains a steady stream of advertising revenue for their commercial sponsors).

Systems that provide support for the social creation of content are sometimes referred to as *social machines* [43]. These systems have been the focus of considerable research attention by the Web Science community [234, 236, 243, 382, 427, 431, 464, 466]. Social machines were originally understood in terms of a division of labor between the social and technological elements of an online system, with human agents performing a creative role and technological elements assigned to more administrative duties [43]. This early approach to understanding social machines has since expanded to include a bewildering array of different views. Social machines have thus been understood as problem-

solving organizations [236, 243], as systems that feature the participatory involvement of social and technological elements in the realization of system-level phenomena [464], as distributed cognitive systems [382], as purely technological systems that are able to engage in social exchanges with other systems [70, 338], and as systems in which human agents participate in computational processes [244, 410].

Given the scope of the present review, this is clearly not the place to undertake a comprehensive survey of social machine research efforts. It is, however, worth highlighting the relevance of social machines from a cognitive science perspective. There are, in fact, a couple of points that are worth mentioning here. The first is that social machines are typically conceived as socio-technical systems, i.e., as systems that feature some form of interactive engagement between a set of social and technological elements. This establishes an obvious point of contact with work in distributed cognitive science [254, 255], especially since the systems of interest in distributed cognitive science are (mostly) systems that involve cognitively-relevant forms of causal commerce between a set of social and technological elements [see 382]. The link with distributed cognition is perhaps all the more obvious if we view social machines as problem-solving organizations [236, 243]. In this case, we can see social machines as being involved in an activity (i.e., problem-solving) that is typically glossed in cognitive terms. The result is that when we view social machines as problem-solving organizations, we are presented with an important opportunity to study social machines as distributed cognitive systems [382].

A second point to note is that issues of complementarity often emerge as a focal point of social machine research. In particular, by virtue of their socio-technical nature, social machines are ideally poised to exploit the respective capabilities of human agents and conventional computational systems [135, 139, 349]. Indeed, it is often the distinctive capabilities of human and machine elements that underlies the interest in social machines as systems with advanced problem-solving capabilities [236]. This emphasis on complementarity opens up important links with work in other areas of computer science, most notably work on complementary computing [269], mixed-initiative computing [245], human-machine

symbiosis [261] and heterotic computing [271]. But in addition to this, issues of complementarity also help to establish a point contact with cognitive science and the philosophy of mind, especially in respect of work into extended cognitive systems (see Section 4). In this case, it is often the complementary nature of a set of biological and non-biological resources that is seen to provide the basis for advanced forms of intelligence [98]. Such claims strike a chord when it comes to the kinds of problems (e.g., disease, world hunger and global climate change) that are typically mentioned in relation to social machines [see 236]. Such problems may be sufficiently taxing as to require the assembly of socio-technical systems that co-opt the complementary computational and representational capabilities of both humans and machines. Indeed, it is perhaps only by merging the properties of human and machine components into some larger functionally-integrated problem-solving ensemble that we will be able to gain sway over problems that lie beyond the computational ambit of either the biological brain or our most advanced forms of technological contrivance.⁶

7.5 Knowledge Engines

Although it is no doubt possible to identify many different kinds of social machine, one kind of social machine is of particular interest. These are social machines that are involved in some form of knowledge-related activity [427, 458], such as the elicitation, acquisition, representation, discovery and exploitation of knowledge. Social machines of this ilk have been referred to as *knowledge machines* [345, 458].

Wikipedia is perhaps the most prominent example of a knowledge machine. It supports a particular form of knowledge acquisition, in which epistemic outputs are generated and refined as a result of collective editing actions (see Section 7.3). The outputs of such processes are, of

⁶Such claims resonate with the views of philosophers like Verdoux [522] who see technologically-mediated cognitive enhancement as a means to circumvent the inherent limitations of the human cognitive system. A similar idea is encountered in respect of the notion of heterotic computing [271]. In this case, socio-technical hybridization is seen to yield a ‘social computer’ whose computational power exceeds the limits of silicon-based digital circuits [244].

course, primarily intended for human consumption, and they thus tend to diverge from the sort of outputs (e.g., computational ontologies) that are delivered by conventional knowledge engineering efforts. Other kinds of knowledge machine, however, participate in the production of more formal knowledge-related resources. These include collaborative ontology authoring environments [438], semantic wikis [285], and a variety of online multiplayer games that provide support for the processes of semantic annotation and ontology alignment [441, 442, 496].

A particularly interesting class of knowledge machine is concerned with the processes of knowledge discovery. For the sake of convenience, we will refer to these knowledge machines as *discovery-oriented knowledge machines*. Discovery-oriented knowledge machines typically rely on large-scale social participation as a means of analyzing bodies of scientific data. Citizen science systems, such as Galaxy Zoo, are one example of this sort of knowledge machine [see 308]. Such systems often contribute to important scientific discoveries. The Galaxy Zoo system, for example, has contributed to the discovery of an astronomical phenomenon known as *Hanny's Voorwerp* [309], as well as a previously unknown class of greenish-colored galaxies, aptly called Green Pea galaxies [76].

Games With A Purpose (GWAPs) are another important category of discovery-oriented knowledge machines. These are systems that attempt to exploit peoples' enthusiasm for playing computer games in order to perform some form of scientifically-relevant information processing [200]. The protein-folding game, Foldit, is a particularly well-studied example of this sort of system [130, 273]. Foldit is an online multiplayer game that aims to derive accurate protein structure models via game-play responses. The game involves the presentation of improperly folded protein structures to human game-players, and the protein structure is then manipulated using a combination of manual and automatic actions so as to maximize the score associated with a computed evaluation metric. The game is interesting because it provides a compelling example of the way in which knowledge machines (and social machines, more generally) can be used to maximally exploit the distinctive (albeit complementary) capabilities of human and machine components [135]. For example, in attempting to maximize their score, an individual user

can interact with the protein structure, tugging and twisting the protein backbone as a means of exploring the target solution space. In the course of these activities, human game-players are deemed to rely on a set of visual and spatial cognitive abilities that are, as yet, unmatched by the capabilities of existing AI algorithms. There is, however, an important role for machine-based processes in supporting the user's search for optimal protein conformations. In particular, the Foldit interface provides access to a range of tools that implement so-called "automatic moves." These include, for example, a "wiggle" routine that attempts to perform a localized search for high-scoring protein structures in the vicinity of the current structural candidate [see 130].

Perhaps one of the most impressive accomplishments of the Foldit system, at least from the perspective of knowledge discovery, relates to its success in deciphering the crystal structure of the retroviral protease of the Mason-Pfizer monkey virus, a simian AIDS-causing virus [274]. The structure of this protein (an enzyme) had remained elusive despite attempts to solve the problem using conventional computational and experimental methods. When assigned to the Foldit system, a group of Foldit players were able to produce an accurate 3D model of the target protein within the space of just three weeks! This represents an important breakthrough for the biomedical research community, especially given the importance of retroviral proteases to Human Immunodeficiency Virus (HIV) research [e.g., 280]. Such results provide us with an important demonstration of the power and potential of GWAPs with respect to their ability to contribute to the process of scientific discovery.

Aside from citizen science systems and GWAPs, it is possible to discern a third type of discovery-oriented knowledge machine. These are systems that attempt to yield insights into a broad array of social phenomena. The general idea is that the Web provides us with an important opportunity to see 'society in action' [cf. 286]. Thus, just as the Web has provided the basis for large-scale forms of social participation in any number of online activities, so too it has also opened the door to novel forms of social observation and analysis [see 485]. Crucially, as our everyday social activities and endeavors become ever-more closely

entwined with the online realm, so it becomes increasingly tempting to see the Web as part of the causally-active physical fabric that *realizes* social processes [see 464]. In other words, the Web presents us with a vision of a networked society in which at least some kinds of social phenomena are subject (at least in part) to Web-based forms of computational realization. This raises a host of important issues concerning our ability to use the Web (or parts thereof) as a form of “digital socioscope” [339] or “social observatory” [87].⁷

As an example of work in this area consider a study by González-Bailón et al. [199], which sought to understand how the distribution of participation thresholds within the Twitter community may have contributed to the emergent profile of the Spanish Indignados protest movement. Here we see a noted social machine—namely, Twitter—being implicated in the forces and factors that work to drive a specific form of collective action. Indeed, González-Bailón et al. [199] suggest that their study helps to illuminate the role of individual communication in shaping patterns of collective behavior:

We use this case to analyse the way in which individual communication patterns concatenate in complex networks of interaction that can ultimately lead to an explosion of activity on the aggregate level, hence boosting the global visibility of social movements and political protests. [199, p. 264]

Using the Web as a platform for social observation and monitoring clearly raises a host of important ethical issues relating to privacy, surveillance and social control. Nevertheless, the ability to observe and monitor social activity in the online realm does provide us with a valuable opportunity to improve our understanding of social processes. This is important, because our contemporary society is a system of such

⁷Interestingly, work in this area often embraces a mechanistic approach to understanding social phenomena—an approach that is in perfect accord with the mechanistic view of social machines alluded to by some social machine researchers [e.g., 464]. Aharony et al. [6], for example, suggest that network-based monitoring techniques can be used to investigate the mechanisms that give rise to social or collective phenomena, in a manner that resembles the use of neuroimaging techniques to help tease apart the mechanistic realizers associated with brain-based cognitive phenomena.

complexity that its dynamical profile often resists our best attempts at prediction and explanation. In the wake of such complexity, it is perhaps tempting to think that the mechanistic underpinnings of social phenomena are doomed to forever lie beyond the reach of our (social) scientific grasp. However, when we see the Web as part of the material fabric of society (i.e., as part of the physical machinery that realizes social phenomena), then we are afforded a much more positive perspective on the empirical and theoretical prospects of contemporary social science. This is because advances in mechanistic understanding (across all the sciences) are often linked to an ability to subject some target system to sophisticated forms of instrumentation and measurement. Perhaps, therefore, we can see the advent of the Web, and the current efflorescence of Web-enabled devices, as marking a potential sea change in our ability to establish an explanatorily- and predictively-potent grip on the social realm. Just as progress in other areas of science has followed hot on the heels of our ability to observe, measure, and monitor—consider the impact of the microscope and telescope on the fields of biology and astronomy—perhaps the Web is poised to progress the cause of the social sciences in a similar manner. In essence, what the Web gives us is an ability to observe (in more-or-less real time) the ebb and flow of social processes on a (potentially) global scale. As a result of such newfound abilities, we may, at last, be able to acquire the sorts of data that informs our search for the mechanistic bases of (at least some kinds of) social phenomena.

It should, of course, be noted that these categories of discovery-oriented knowledge machine (i.e., citizen science systems, GWAPs and social observatories) are not mutually exclusive. There is, as such, no reason why a system such as Foldit could not be regarded as *both* a GWAP and a citizen science system. Likewise, there may very well be systems that qualify as members of all three categories of discovery-oriented knowledge machine—social science experiments using virtual worlds, such as Second Life or Pokémon Go, may be a case in point [see 22].

It is also important to note that each kind of discovery-oriented knowledge machine may be amenable to further forms of conceptual decomposition. When it comes to GWAPs, for example, Smart [458]

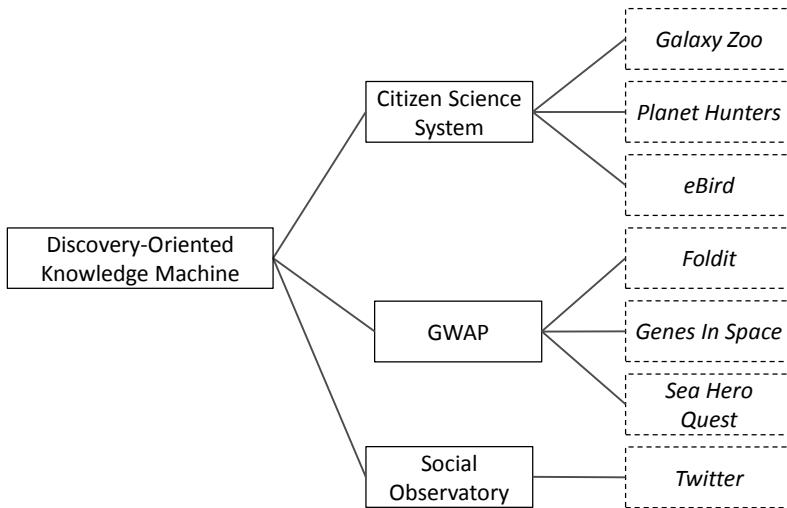


Figure 7.2: A taxonomy of knowledge machines, with a particular focus on discovery-oriented knowledge machines. The knowledge machine instances referred to in this figure (symbolized with dashed lines and italicized font) include Galaxy Zoo [310], Planet Hunters [425], eBird [488], Foldit [130], Genes In Space [118], Sea Hero Quest [359, 473], and Twitter [199].

suggests that we should recognize at least two kinds of GWAP, namely, goal-transparent and goal-opaque GWAPs. The difference between these two subcategories is said to be grounded in the relative ‘visibility’ of the game’s epistemic objectives to the end-user. In the case of Foldit, for example, the link between game-play responses and the epistemic (or scientific) objectives of the game are easily accessible to members of the game-playing community. Such is not the case, however, for goal-opaque GWAPs. In this case, the game-playing community may not even be aware that their game-play responses are being used for the purposes of scientific analysis.

Finally, it is important to note that many other categories of knowledge machine could be identified in future work. The focus of the present section has been on one particular kind of knowledge machine—i.e., discovery-oriented knowledge machines (see Figure 7.2)—but this does not necessarily reveal the true extent of the knowledge machine con-

cept. In addition to discovery-oriented knowledge machines, future work is likely to reveal knowledge machines that support other kinds of knowledge-oriented process, such as those associated with the acquisition and representation of existing (i.e., pre-discovered) knowledge. The result is likely to be an extended taxonomy that highlights the scope and richness of the knowledge machine concept. Given that such forms of conceptual analysis fall within the remit of conventional knowledge engineering [429], we can easily envisage a state-of-affairs in which the process of deriving an extended taxonomy of knowledge machines is itself something that is performed by a knowledge machine [see 466].

7.6 Leaning on the (Social) Environment

In Section 7.1, we saw that issues of network structure are relevant to our understanding of collective cognitive organizations. In particular, network structure emerged as one of the factors that influenced the cognitive performance of a group, team or some other collective. This highlights one of the major goals of contemporary network science: to identify those network-related properties that provide us with a predictive or explanatory grasp of social (i.e., system-level) phenomena. However, it is not just the properties of the larger networked ensemble that are of interest to network scientists. Network scientists also seek to understand how structural or topological properties affect the behavior of the network's nodal elements (e.g., human individuals). To help illustrate this, consider a classic ethnographic study by Elizabeth Bott [62]. Bott examined 20 urban British families and attempted to explain the considerable variation in the way husbands and wives performed their family roles. In some families, there was a strict division of labor: husband and wife carried out distinct household chores separately and independently. In other families, the husband and wife shared many of the same tasks and interacted as equals. Bott found that the degree of segregation in role-related behavior varied with the connectedness (or density) of the family's social network: the more connected the network, the more likely the couple were to maintain a traditional segregation of husband and wife roles. This showed that the structure of the larger network in which the couple were embedded played an important role

in influencing the couple's behavior. From an explanatory perspective, what we encounter here is a shift in explanatory focus. In particular, by focusing on the structure of a social network we are adopting something of an 'externalist' perspective with respect to the explanation of individual (agent-level) properties. This is important, because it might be thought that the best place to look for factors that explain or predict the properties of individual agents is at the level of the individual agents themselves. Thus, if we wanted to predict whether a particular individual was likely to be depressed or obese, we might be inclined to focus our attention on the characteristics of the individual rather than the features of the environment in which that individual was situated. And yet, the results of recent work in network science suggests that the structure of the social environment, specifically the relationships between individuals, may be of explanatory relevance with respect to both depression [411] and obesity [94] (see Christakis and Fowler [95] for a review).

When it comes to understanding the properties of individual's agents, therefore, network science encourages a shift in attention—away from the individual agent and towards aspects of the agent's (social) environment. This point is nicely summarized by Borgatti et al. [61]:

Whereas traditional social research explained an individual's outcomes or characteristics as a function of other characteristics of the same individual (e.g., income as a function of education and gender), social network researchers look to the individual's social environment for explanations, whether through influence processes (e.g., individuals adopting their friends' occupational choices) or leveraging processes (e.g., an individual can get certain things done because of the connections she has to powerful others). [61, p. 894]

Interestingly, this shift in explanatory focus—from the individual to the environment—resembles the sort of shift encountered in the case of extended cognition (see Section 4). In fact, the only substantive difference is that extended cognition is concerned with cognitive phenomena, while network science (at least in a social science context) restricts

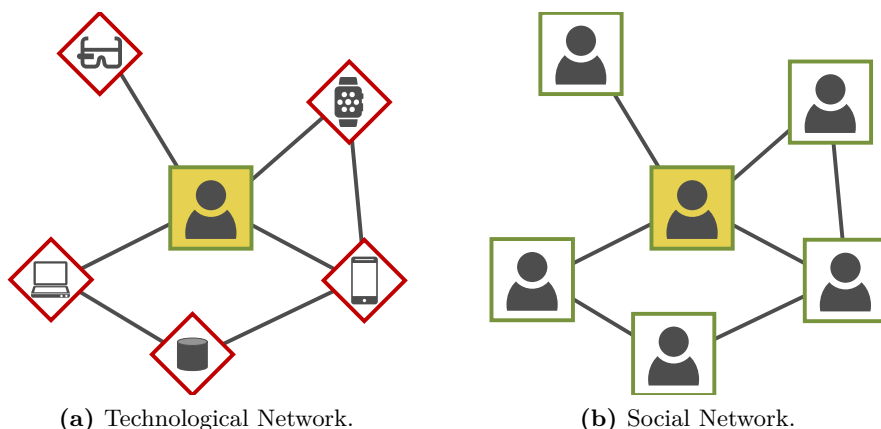


Figure 7.3: Network-oriented analyses of agent-level properties with respect to (a) technological resources and (b) other agents. In both cases, the properties of the focal agent (shaded icon) are explained (at least in part) by referencing the relationships established with extra-agential elements.

its empirical focus to social relationships. In other words, extended cognition imposes a constraint on the nature of the explanandum (i.e., cognitive phenomena), whereas social network analysis imposes a constraint on the explanans (i.e., social network structure). Apart from this difference, however, there is an obvious commonality between the two approaches in terms of the appeal to explanatorily-significant factors that lie beyond the traditional biological borders of the individual agent. If, for example, we broadened the remit of social network analysis to include the linkages with artefactual (as opposed to purely social) elements, then it seems perfectly possible that we would seek to explain the cognitive properties of the individual agent with respect to the time-variant structural profile of informational circuits that connect the biological brain with a surrounding nexus of extra-agential (artefactual) resources (see Figure 7.3). This helps us see the theoretical and empirical relevance of network science to our understanding of extended cognitive systems [see 459].

This way of thinking about the relationship between extended cognition and social network analysis is also something that is relevant to

debates about a specific kind of extended cognition, namely, *socially-extended cognition* [182, 505]. For the most part, philosophical treatments of extended cognition tend to limit their scope to systems that involve bio-technological forms of bonding. Socially-extended cognition, however, extends the scope of extended cognition to include the social environment. In this case, the main focus of attention relates to whether the elements of the social environment (e.g., other human agents) are of constitutive relevance to the realization of an individual's cognitive processing routines. From a network-theoretic perspective, socially-extended cognition requires us to focus our attention on the informational circuits that connect a focal individual with other individuals. Inasmuch as we see such agent-involving circuits as forming part of an extended cognitive mechanism that realizes the cognitive states and processes of a particular individual, then we have a form of cognition that is *both* a form of extended cognition (i.e., socially-extended cognition) and a form of collective cognition (in the sense that multiple individuals are component elements of an extended cognitive mechanism).

The possibility of socially-extended cognition is recognized by at least some proponents of extended cognition. In their original treatment of the extended mind thesis, for example, Clark and Chalmers [108] suggest that there is no reason, at least in principle, why socially-extended forms of cognition should not exist:

Could my mental states be partly constituted by the states of other thinkers? We see no reason why not, in principle. In an unusually interdependent couple, it is entirely possible that one partner's beliefs will play the same sort of role for the other as the notebook plays for Otto. [108, p. 17]

In spite of all this, claims about socially-extended cognition remain controversial. This is due, at least in part, to worries about the extent to which other agents can meet the sort of criteria required for cognitive extension, for example, the trust-and-glue criteria proposed by Clark and Chalmers [108].

It is here that a focus on the Web, and especially the Social Web, could be philosophical value. This is because the Web plays a crucial

role in shaping the nature of our social interactions and engagements, transforming the way in which a broad array of social activities are undertaken. As a result of this transformation, it is possible that the Web may help to surmount (via technological means) some of the barriers that would otherwise prohibit the practical realization of socially-extended forms of cognition. Thus, while conventional forms of social interaction are clearly limited by a range of practical issues and concerns (e.g., the need for co-location), it is possible that Web technologies are able to alleviate some of these constraints. We can thus imagine Web technologies being specifically designed so as to yield conditions that are conducive to socially-extended cognition. Such conditions may be difficult to achieve in the case of conventional face-to-face interactions, and it is for this reason that the (Social) Web can be seen to provide a new impetus to debates regarding the practical possibility of Web-based forms of socially-extended cognition [see 454].

Some insight into the opportunities the Web provides for socially-extended cognition comes from applications that rely on real-time or continuous crowdsourcing [see 299]. In general, these applications seek to provide cognitive support to a particular individual by enlisting the help of other individuals. The VizWiz system, for example, is a system that seeks to support blind people in dealing with the challenges of a visual environment [50]. VizWiz enables blind individuals to upload images from a smartphone and then receive descriptions of the image from other individuals in what is described as “nearly real-time.” Such systems provide a clear and compelling demonstration of the idea that the wider social environment can be recruited into a form of cognitively-potent information processing—one that influences the thoughts and actions of the visually-impaired individual. There is an obvious parallel, here, with the classic case of Otto and his notebook in the conventional extended mind case [108]. Despite some clear differences in the nature of the cognitive processing routines that are being performed (e.g., visual processing versus mnemonic recall), as well as differences in the nature of the relevant material realizers (i.e., socio-technical system versus notebook), the two cases are roughly equivalent: in both cases, we have a form of disability that is being addressed by virtue of the kind of

engagements that are made with respect to the extra-organismic (in this case, socio-technical) environment.

Another case of real-time crowdsourcing comes in the form of the Soylent system [48]. Soylent is a word processing technology that aims to use crowdsourcing techniques as a means of supporting the task of document writing and editing. In particular, Soylent is described as:

[A] word processing interface that utilizes crowd contributions to aid complex writing tasks ranging from error prevention and paragraph shortening to automation of tasks such as citation searches and tense changes. [48, p. 85]

Such forms of socially-distributed document authoring establish an interesting point of contact with debates about the extended mind. This is because appeals to the role of extra-organismic resources in the creation of written artefacts are a common focus of attention for the extended mind community. Consider, for example, the claim that writing can (on occasion) be regarded as a form of extended thinking. In the introduction to his book, *Supersizing the Mind*, Clark [102] reflects on a discussion between Richard Feynman and Charles Weiner concerning the role of writing in Feynman's intellectual activity. Clark proposes that:

... Feynman was actually *thinking* on the paper. The loop through pen and paper is part of the physical machinery responsible for the shape of the flow of thoughts and ideas that we take, nonetheless, to be distinctively those of Richard Feynman. [102, p. xxv, original emphasis]

Needless to say, the kind of interactions that users have with collaborative authoring systems like Soylent are probably unlike the bio-artefactual exchanges associated with more insular forms of writing activity. Nevertheless, in accounting for why it is that we are able to generate a written text using Soylent, it may be that the contributions of others are seen to play an explanatorily-crucial role—one that is perhaps functionally akin to the sort of role played by extra-organismic resources in shaping the profile of our more solitary orthographic endeavors [see 96, pp. 206–207]. The key point, here, is not that Soylent serves as the

de facto technological substrate for a specific form of socially-extended cognizing. Rather, the point is that, in view of the present discussion, we can perhaps begin to see how research in Web Science is able to establish a productive point of contact with contemporary philosophical debates regarding the role of extra-organismic resources in extending the reach of the human cognitive system.

7.7 Mandevillian Intelligence

When it comes to issues of collective cognition and collective intelligence, a common focus of research attention relates to the forces and factors that determine the quality of cognitive outputs. In the case of collective sensemaking, for example, it is common to direct attention to those features that enable a group of sensemaking agents to interpret situation-relevant information in ways that line up with facts about reality. Such features include the properties of individual agents (e.g., their susceptibility to confirmation bias), the patterns of communication established by the agents, and the properties of the technological system that works to coordinate collective efforts. In many cases, the properties of the technological system are of particular interest, for these are the properties that can be tweaked as part of the design and development process. In attempting to optimize the performance of (e.g.) a collective sensemaking system, it thus makes sense to design the technological components in such a way as to ensure that the larger systemic organization achieves the best sensemaking outcome, perhaps by regulating the flow of information between individual sensemakers, or by attempting to mitigate the presence of individual forms of cognitive bias [see [124](#)].

It is at this point that we encounter a novel proposal concerning the cognitive properties of individual agents and the quality of the outcomes that are delivered as a result of collective cognitive processing. According to this proposal:

Cognitive and epistemic properties that are typically seen as shortcomings, limitations or biases at the individual level can, on occasion, play a positive functional role in support-

ing the emergence of intelligent behavior at the collective level. [455, p. 3]

In other words, the claim is that the seemingly sub-optimal cognitive performance of individual agents can, on occasion, play a productive role in helping to bring about desirable outcomes at the collective level of analysis. The form of intelligence that emerges in these situations is called *mandevillian intelligence* [455, 456]. Mandevillian intelligence is thus a specific form of collective intelligence, in which intelligent outcomes at the collective level are seen to arise from the presence of cognitive properties that we generally associate with poor performance at the level of individual agents (e.g., poor mnemonic recall, defective reasoning, limited attentional capacity, and so on).

Although mandevillian intelligence is no doubt contentious, it is an idea that is of considerable importance to cognitive science, philosophy,⁸ and (as we shall see) Web Science. From the standpoint of cognitive science, the notion of mandevillian intelligence encourages us to rethink our traditional approach to the design of collective cognitive systems. If, for example, we accept that individual cognitive limitations can sometimes play a productive role in yielding collective forms of cognitive success, then it behooves the designers of collaborative technology to question the logic of design strategies that always attempt to eliminate (or at least attenuate) such limitations. Indeed, it may very well be the case that in our attempts to maximize collective cognitive performance, we should consider design strategies to that seek to degrade and diminish the cognitive capabilities of the individual [455].

Such claims obviously raise a host of ethical issues concerning the extent to which the cognitive competencies of the one (the individual) should be traded against the cognitive capabilities of the many (the collective), and few will probably be willing to embrace an approach to technology design that sacrifices the capabilities of the individual agent for the sake of the collective cognitive good. But even if we retreat from the idea of diminishing individual cognitive capabilities

⁸The philosophical implications of mandevillian intelligence are perhaps most keenly felt in the case of contemporary epistemology, especially virtue epistemology [see 456].

via the introduction of new technologies, we can at least begin to see how the notion of mandevillian intelligence alters our view as to the cognitive and epistemic value of a particular technology. A technology that degrades the cognitive capabilities of a given individual may thus be subject to negative evaluation, and rightly so. But from the standpoint of mandevillian intelligence, such evaluations are specific to individual cognition. Once the focus of our analytic gaze shifts from the individual to the collective level—from individual cognition to collective cognition—we might be able to see the technology in a somewhat different light. The technology may, in fact, be subject to positive evaluation on the grounds that it exploits an individual cognitive shortcoming for the purposes of enhancing the performance profile of a collective cognitive system.

In a Web Science context, the significance of mandevillian intelligence starts to come into clearer view once we reflect on claims about the cognitive impacts of the Web. Such impacts have seldom been viewed in a positive light. Carr [80, 81], for instance, has expressed concerns about the deleterious effects of the Web and Internet on aspects of (individual) cognitive performance. He worries that the nature of our current engagement with the Web is leading to a general decline in mnemonic and attentional functioning. Similar concerns are sometimes expressed in relation to specific Web technologies, such as personalized search engines [351, 439]. The worry, in this case, is that by delivering a set of personalized (and therefore filtered) search results, personalized search engines have the potential to undermine epistemic standing by accentuating existing forms of cognitive bias. The result, in most cases, is an appeal to interventions that are intended to ameliorate the negative epistemic sequelae of personalized search technology. These include changes in user behavior and corporate policy, backed up by government regulation [439].

It is here that the notion of mandevillian intelligence helps to reshape (or at least refocus) the nature of the ongoing debate. For even if we accept that personalized search is, in general, injurious to an individual's epistemic health, this does not mean that personalized search is bereft of *any* sort of epistemic benefit. In particular, it is far from clear that the epistemic consequences of personalized search for a community

of Web users is exactly the same as for the individual members of that community. From the perspective of mandevillian intelligence, for example, we might view personalized search as playing a productive role in maintaining a degree of cognitive diversity within a community of Web users. Given that such diversity is sometimes seen to play an important role in enabling a collection of human individuals to discover, resolve or otherwise track the truth in some domain of interest [341, 563],⁹ it seems that personalized search may, in fact, help to support the epistemic functioning of a community, even if the same technology also acts to undermine the epistemic credentials of the individual agent. In essence, what the notion of mandevillian intelligence gives us is a means of avoiding a rush towards premature judgements regarding the cognitive or epistemic value of the Web, or a particular component thereof. Just because a particular technology, such as personalized search, turns out to have little or no benefit at the level of individual agents, this does not mean that it has no benefits *tout court*. This is perhaps of particular importance when one considers the recent criticisms that have been leveled at major technology providers (e.g., Facebook and Google) by a number of prominent political leaders.¹⁰ The main value of mandevillian intelligence, in this respect, is that it helps to inject an additional epistemological dimension into the ongoing political debate. In particular, it is unclear whether the criticisms of major technology vendors are justified in the absence of a better understanding of how (e.g.) information filtering algorithms affect the cognitive performance of teams, groups, communities and societies—and not just the performance of human individuals. Only in the light of this enhanced (and empirically-informed) understanding will we be in a position to make judgements concerning the cognitive and epistemic consequences of the Web and the particular kinds of interventions (e.g., government regulation) that might be required to ensure its virtuous (i.e., cognition-friendly) operation.

⁹In particular, by limiting attention to specific subsets of information based on (e.g.) a user's browsing history, personalized search may work to diversify access to information and thus prevent premature forms of cognitive convergence.

¹⁰See <http://www.bbc.co.uk/news/technology-37798762>.

8

Extending the Butterfly

Web Science is a highly interdisciplinary endeavor, with well-established links to physics, mathematics, computer science, AI, sociology, psychology and other disciplines [235, 428]. As we saw in Section 1, the interdisciplinary nature of Web Science is nicely captured by the Web Science Butterfly Diagram (see Figure 1.1), which highlights some of the topics belonging to the Web Science research agenda. In its original form, however, the Butterfly Diagram does not include the disciplines that have been the focus of the present review, i.e., philosophy of mind and cognitive science. In the wake of the present review, this seems like an important omission. Although some of the topics highlighted by the present review could perhaps be subsumed within the existing disciplines of psychology, sociology, and artificial intelligence, it should by now be clear that the Web raises a host of issues that dovetail specifically with the interests of cognitive science and the philosophy of mind. The present review helps to reveal at least some of these issues; however, it is by no means an exhaustive survey of this part of the intellectual terrain. The cognitive properties of the Web (or at least the points of cognitive scientific interest for Web Science) include an array of issues that were not part of the present review (e.g., the implications of the

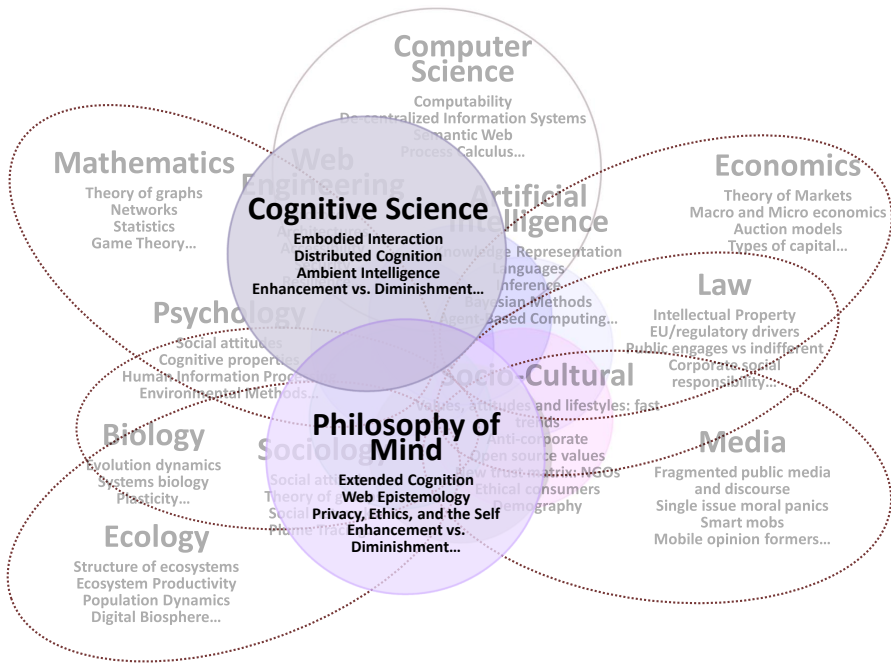


Figure 8.1: The Web Science Butterfly Diagram (extended).

Web for machine-based cognitive capabilities [451, 452]). In addition, some of the issues that did form of the part of the present review could have benefited from a more detailed discussion. The work relating to memory, for example, is of sufficient complexity as to warrant separate treatment [71, 110, 113, 114, 232, 281, 335, 337, 367, 445, 472, 531].

The upshot of all this is a proposal to ‘extend the butterfly’—to incorporate cognitive science and the philosophy of mind into an extended version of the Web Science Butterfly Diagram (see Figure 8.1). The primary aim of such extensions is, of course, to highlight the points of interdisciplinary contact between Web Science, cognitive science, and the philosophy of mind. In addition to this, however, the extensions also tell us something about the transformational nature of the Web and the metamorphic character of the science that studies it. The Web is not a static system with immutable properties; rather, it is a constantly changing, dynamic system. In view of the Web’s protean nature, it is

perhaps incumbent on all disciplines to periodically assess the relevance of the Web to their own areas of intellectual study. Cognitive science and philosophy of mind are, of course, no exception. Not only do such disciplines bring a range of insights, ideas, methods, and techniques that are relevant to Web Science, they are also disciplines whose own research agendas are likely to be influenced by the shifting technological and social landscape of the Web.

The aim of the present section is to reinforce the case for an expansion of the interdisciplinary borders of Web Science. In other words, we aim to highlight some of the reasons why the wings of the butterfly should be extended to incorporate the disciplines of philosophy and cognitive science. Our approach is to reach across the various issues and topics that have been discussed in previous sections and identify some of the more prominent cross-cutting themes. In doing so, we hope to make the case for our particular form of ‘lepidopteran’ expansionism. But in addition to this, we also hope to reveal a number of areas that could serve as the basis for future work.

8.1 Cognitive Connections

Although it is probably fair to say that cognitive science and the philosophy of mind are seldom seen to lie at the heart of the Web Science research agenda, it should by now be clear that a number of important linkages exist between these disciplines. As a means of bringing these linkages into sharper focus, Table 8.1 identifies a number of traditional Web Science topics that were discussed as part of this review.

In addition to highlighting the nature of the interdisciplinary connections between Web Science, cognitive science and the philosophy of mind, the present review also demonstrated the use of the Web to further our understanding of the human cognitive system. In Section 3.6, for example, we encountered the idea of the Web being used as an experimental or observational platform—one that could be used to shed light on the features of both individual and collective cognition [197, 327, 473].

We also sounded a cautionary note about claims concerning the cognitive impacts of the Web (see Sections 4.1 and 4.3). In particular, we suggested that our interpretation of research results can be

Table 8.1: Web Science topics discussed as part of the present review.

Web Science Topic	Cognitive Relevance	Sections
Semantic Web	Extended Cognition	4.1
Linked Data	Extended Cognition	4.1
Trust	Extended Cognition	4.2, 4.6
Privacy	Extended Cognition	4.7
Social Web	Social Cognition, Collective Cognition	6, 7
Human Computation	Collective Cognition	7.5
Social Machines	Collective Cognition	7.4
Quantified Self	Embodied Cognition	3.5, 3.6
Personalization	Embedded Cognition, Extended Cognition	5.3, 4.2
Ambient Intelligence	Embedded Cognition	5.2, 5.3
Wearable Devices	Extended Cognition	3.4
Social Bots	Social Cognition	6.2

influenced by a number of factors, and these often limit the validity of blanket statements to the effect that the Web is (e.g.) degrading our cognitive capabilities. The following points are thus intended to serve as a safeguard against what might be called the *sin of over-generalization*:

- **Task Specificity:** Research results can be influenced by the nature of the task that is being performed by Web users. In the case of neuroimaging studies, for example, the sorts of results obtained with Web search tasks [443] are not necessarily the same as those obtained in the context of socially-oriented tasks [343].
- **Cognitive Function Specificity:** The effects obtained in respect of one particular kind of cognitive processing (e.g., memory) should not be generalized to other kinds of cognitive processing (e.g., problem-solving). This also applies to the distinctions between cognitive functions that are of the same type (e.g., different forms of memory). For example, the cognitive implications of the

Web may vary according to whether our research interests relate to semantic memory or autobiographical memory.

- **Cognitive System Specificity:** We should recognize that the cognitive gains/losses enjoyed/endured by individual cognizers are not the same as those encountered in the case of collective cognition (see Section 7.7). The nature of the focal cognitive system (e.g., collective cognitive system or individual cognitive agent) thus affects our interpretation of research results (see Section 8.3 for more on this).
- **Access Specificity:** The ways in which we access the Web are likely to affect the nature of research results. We should not assume that the nature of our interactive engagements with the Web are exhausted by the use of one particular kind of technology (e.g., search engines), or that HTML-based modes of information presentation limit the nature of our cognitive contact with the online environment.
- **User Specificity:** Different categories of users may respond to the Web in different ways. For example, the cognitive effects observed with “digital natives” may not be the same as those observed with “digital immigrants.”
- **Historical Specificity:** The Web is a highly dynamic system, forever apt to change in response to technological innovation and changes in social practice. We should not, therefore, see the properties of the *current* Web as limiting the scope of its future cognitive impact. To help us see this, we need only reflect on the way in which our interactions with the Web have changed over the course of the Web’s (rather brief) history.

We suggest that Web Scientists, cognitive scientists and philosophers would do well to consider these forms of specificity in the context of future research efforts.

8.2 Boundary Dissolution

Entanglement, merger and incorporation are themes that can be found throughout the Web Science literature. Such themes are especially prominent in the more sociologically-oriented areas of Web Science. The Web and Internet are thus sometimes seen to contribute to an important form of boundary dissolution that blurs the traditional distinction between the ‘social’ and the ‘technological’ [316, 317]. “For some theorists,” writes Lupton [317], “the very idea of ‘culture’ or ‘society’ cannot now be fully understood without the recognition that computer software and hardware devices not only underpin but actively constitute selfhood, embodiment, social life, social relations and social institutions” (p. 2). A similar view is expressed by Smart and Shadbolt [464]. They suggest that the Web and society have become inextricably linked, to the point where the Web serves as part of the physical machinery (the mechanistic substrate) of the social world. The Web, in this sense, emerges as part of the material infrastructure that (at least in part) makes our contemporary society what it is.

In addition to the forms of boundary dissolution occurring in respect of the social and the technological, the Web and Internet are also seen to dissolve the boundaries that exist between a number of other previously distinct concepts. Notable examples include the distinctions between *work* and *leisure* [183], between the *public* and the *private* [503], between the *physical* and the *virtual* [317], between the *human* and the *machine* [315, 521, 532], and between the *online* and the *offline* [175, 177].¹

The present review expands on these forms of boundary dissolution. One example of this comes from the discussion of corporeal extension (see Section 3.3), which highlighted the forms of entanglement, merger and incorporation that might occur between the biological body and Web-enabled devices. Such ideas are consistent with views that have been developed by both cognitive scientists and social scientists. In a

¹O’Hara (personal communication, 10 October, 2013) notes that the boundary between cybercriminality and libertarianism is also somewhat indistinct, with Distributed Denial of Service (DDoS) attacks on corporate websites being viewed as a crime by some and a legitimate form of civil disobedience by others.

cognitive science context, for example, Gibbs [186] argues that the more technological devices extend the body and integrate non-organic material with human flesh, the more they “work to dissolve any clear boundary between bodies and world” (p. 18). Similar views are expressed by the sociologist Deborah Lupton [315]. She notes that wearable technologies not only “become prosthetics of the body but extend the body into a network with other bodies and objects” (p. 27).

An additional form of boundary dissolution can be found in respect of the notion of extended cognition (see Section 4). The main focus of attention, in this case, is the extent to which the Web supports a form of bio-technological hybridization—one in which the informational and technological elements of the online environment can be seen as literal constituents of the physical fabric that realizes human cognitive states and processes. Such forms of integration and merger are sometimes seen to blur the boundary between the cognitive agent and the agent’s ‘local’ technological environment [99, 102]. The traditional biological borders of the cognitive machine (i.e., the human brain) may thus be breached as we begin to appreciate the way in which our cognitive profiles are shaped by the forms of contact we have with the Web. As Clark [99] notes:

As we move towards an era of wearable computing and ubiquitous information access, the robust, reliable information fields to which our brains delicately adapt their routines will become increasingly dense and powerful, further blurring the distinction between the cognitive agent and her best tools, props and artifacts. (p. 275)

References to the notion of extended cognition can be found throughout the present review. In addition to the primary discussion in Section 4, issues of cognitive extension also emerged in relation to embodied cognition (see Section 3.5), social cognition (see Section 6.4), and collective cognition (see Section 7.6). These discussions drew attention to the potentially hybrid character of human cognitive capabilities—the fact that many forms of cognitive accomplishment may originate from information processing routines that straddle both the biological and

technological domains. One of the points to emerge from the discussion of cognitive extension was the mutability of human cognition—the idea that the cognitive capabilities of the cognitive agent are seldom fixed by the processing capacities of the bare biological brain. The curves of the human cognitive system, if this is correct, can perhaps never be fully resolved simply by virtue of our attempts to trace the convoluted contours of the cerebral cortex. The notion of extended cognition thus encourages us to see cognitive phenomena as tied to the operation of environmentally-extended circuits, i.e., circuits that extend beyond the brain to encompass the informational and technological elements of the online, Web-based realm. The upshot of all this is that the human cognitive system has something of a chameleonic character: by virtue of our capacity for cognitive extension, we are able to extend the reach of our cognitive capabilities, and in doing so, the profile of human cognition is itself transformed. It is this capacity for mental metamorphosis that is perhaps one of the defining features of human intelligence [see 98]. And yet, in the process of morphing from one kind of cognitive critter to another, the inherent lability of human intelligence is itself revealed.

All of this helps to introduce us to a further form of boundary dissolution—one that relates to our more general views of intelligence. In Section 7, for example, we saw that the social and collaborative properties of the Web help to lay the foundation for various forms of collective cognition, including socially-extended cognition (see Section 7.6). Such forms of collective cognizing challenge the erstwhile crisp distinction the intelligence of the one (i.e., the human individual) and the intelligence of the many: in leaning on the social environment, we incorporate the cognitive capabilities of others into the material fabric that makes ‘individual’ forms of intelligence materially possible.

Much the same could be said about the distinction between human and machine intelligence. Although the present review did not aim to assess the relevance of the Web from the perspective of machine cognition, it is important to note that all the forms of cognition we have discussed in the present review are applicable to cognitive systems of the computational or technological kind [451, 452, 454, 463]. Of particular interest, in this respect, is the notion of Human-Extended Machine

Cognition (HEMC), which sees Web-based forms of socio-technical entanglement as playing a crucial role in the realization of future forms of ‘machine’ intelligence [452, 454, 457]. Relative to such a view, it is, at best, unclear as to whether we should draw a neat dividing line between natural and synthetic forms of intelligent system—between the cognitive capabilities of the ‘human’ and the cognitive capabilities of the ‘machine’.

There is thus a growing sense that claims about the extended character of human cognition blur the distinctions between a number of forms of intelligence. Once familiar dichotomies (e.g., human intelligence vs. machine intelligence, natural intelligence vs. artificial intelligence, and individual intelligence vs. collective intelligence) are called into question as soon as we begin to appreciate the increasingly bio-technologically hybrid and socially-distributed nature of our intellectual, scientific and problem-solving accomplishments.

Having said all this, it is important to realize that cognitively-potent forms of entanglement, merger and incorporation are not a given in the future network society. As with other forms of boundary dissolution, much depends on the properties of Web technology, as well as the social responses elicited by those properties. As Norman Augustine’s [18] *Second Law of Socioscience* puts it: “For every scientific (or engineering) action, there is an equal and opposite social reaction.” Web-based forms of bio-technological bonding provide us with opportunities for cognitive and epistemic transformation, but they also present us with a number of risks and challenges (recall the discussion concerning the privacy implications of the Web-extended mind in Section 4.7). Concerns over privacy, trust, and security may result in individuals withdrawing from the Web, leading to the re-establishment of traditional borders and boundaries.

8.3 Augmented Cognition?

A recent focus of debate within the Web Science community relates to the impact of the Web and Internet on human cognitive capabilities [81, 312, 333, 448]. For the most part, the tenor of this debate has been largely negative, with key protagonists, such as Nicholas Carr [80, 81], arguing

that the Internet is undermining our ability to think, read and remember. Instead of enhancing our ability to concentrate, Carr argues, the Internet is undermining our capacity for sustained attention, as well as our ability to think deeply about a topic. The result of all this, according to Carr, is a curtailment and fragmentation of otherwise temporally-protracted episodes of cognitive activity, coupled with the adoption of highly superficial forms of cognitively-relevant information processing.

In the wake of the present review, it should be clear that claims of this sort are in danger of committing the aforementioned sin of over-generalization (see Section 8.1). In particular, when it comes to the Web, it is important to remember that any cognitive effects are likely to depend on our modes of access to the Web, and these are likely to change as the result of technological innovation. A range of emerging Web technologies, as well as changes in the way we use the Web, all contribute to an ever-changing landscape against which our notions of the Web's cognitive impact are always likely to be somewhat ephemeral. We should, as a consequence, be wary of blanket statements to the effect that the Web is undermining our cognitive capabilities.

In addition to these objections, it should by now be clear that appeals to the generic notion of *cognition* are themselves a potential source of confusion when it comes to claims about the cognitive impact of the Web. It is crucial, therefore, that in talking about the cognitive sequelae of the Web, we are explicit about what we mean by the term "cognition." By this, we do not mean that it is important to stipulate the kinds of cognitive processes one is interested in (e.g., memory, attention, thinking, problem-solving, and so on); rather, we need to be mindful of the distinction between different forms of cognition (e.g., the distinction between extended, embedded and collective cognition). In discussing the cognitive impacts of the Web, Carr and others focus their attention at the level of individual human agents, and for the most part, their worries and concerns relate to brain-based forms of cognitive processing. The problem with this approach is that it ignores the variegated nature of the cognitive systems that may be brought into existence on the back of our interactions with the Web. In the section on collective cognition, for example, we saw how the Web may yield collective cognitive systems

whose cognitive capabilities are not the same as those exhibited by the individual agents that make up those systems (see Section 7). In view of this, we should ask ourselves whether it is particularly prudent to overlook or disregard the purported benefits of the Web when it comes to issues of collective cognition and collective intelligence. Similarly, why should the notion of *cognition* be limited solely to the realms of human intelligence? Perhaps the cognitive significance of the Web should be judged relative to its implications for machine intelligence, in addition to its purported effects on the human cognitive system [451, 452].

Much the same could be said about the putative role of the Web in supporting the emergence of extended cognitive systems (see Section 4). Indeed, the notion of extended cognition has often been seen as a fitting antidote to the alarmist rhetoric sounded by Carr [81], Greenfield [206], and others. Friends of extended cognition have thus sought to defuse claims about cognitive diminishment by appealing to the cognitive capabilities of the larger systemic organizations (e.g., Web-extended minds) that emerge from human–Web interactions [444]. The key point, here, is that the cognitive capabilities of the larger extended ensemble are not the same as those of the biologically-bounded individual. Thus, even if it could be shown that the Web was undermining the cognitive capabilities of the bare biological brain, it does not follow that the cognitive capabilities of the technologically-extended human individual are similarly diminished. The result, perhaps, is that we should not attempt to vilify the Web based on effects that pertain solely to the bio-cognitive system: changes in neurocognitive function—as revealed by (e.g.) neuroimaging studies [443]—do not necessarily impugn the status of the Web as a cognitively-valuable resource.

8.4 Web Epistemology

The Web is emerging as an important focus of epistemological attention. As is evidenced by the current preoccupation with “false facts” and “fake news,” there is widespread concern (especially in the popular press) about the potential of the Web to deceive and misinform [see 390]. Such concerns have become particularly apparent in the wake of the 2016 U.S. presidential election, where issues of truth and the

factive status of online information became a prominent focus of media attention.² Such issues and concerns dovetail with the interests of contemporary epistemology, especially with the epistemological sub-disciplines of social [195] and applied epistemology [88, 117, 465]. The epistemic implications of the Web have, of course, been the target of previous epistemological research [117, 166, 167, 194]; however, the potential of the Web to inform (and misinform) arguably calls for a much more directed epistemological effort, albeit one that co-opts the empirical and technical expertise of the Web Science community.

The epistemological significance of the Web is highlighted by the present review. Indeed, issues of epistemological significance can be discerned for each of the forms of cognition discussed in previous sections. These issues include the following:

- **Self Knowledge (Embodied Cognition):** Issues of self-related knowledge surfaced in Section 3.5. We saw that the quantified-self movement can be conceived as an epistemological exercise that involves the tracking of personal information. The reliability of the techniques and technologies associated with self-tracking may contribute to improvements in self knowledge. However, in the process of quantifying the self, a digital self may be created. The emergence of this digital self, and the way it interacts with the original self, potentially complicates our understanding of what it means to know the ‘self’ [see 191].
- **Extended Knowledge (Extended Cognition):** Section 4.5 introduced the notion of extended knowledge, which has been the focus of recent epistemological attention [107, 379, 393, 453]. The extent to which we can think of the Web as effectively supersizing our epistemic capabilities [see 314] was discussed in Section 4.6.
- **Situation Awareness (Embedded Cognition):** In Section 5.3, we encountered the idea that Web-enabled devices, coupled with the use of rule-based ‘programming’ techniques (e.g., IFTTT), could be used to enhance an individual’s awareness of (possibly re-

²<http://www.bbc.co.uk/news/world-38168792>

mote) situations, events and contingencies. This establishes a natural point of contact with issues of situation awareness, knowledge and understanding in the human factors community [23, 164, 462]. Such links may be of potential value in helping us understand the notion of situation awareness from an epistemological perspective.

- **Interpersonal Epistemology (Social Cognition):** Interpersonal epistemology was discussed in Section 6.1. Interpersonal epistemology concerns the knowledge that we have of others via the social psychological processes of person perception and impression formation [40]. In a Web-based context, such knowledge derives from the properties of the online systems (e.g., social networking sites, dating sites, etc.) that we use to interact with other individuals.
- **Knowledge Engines (Collective Cognition):** Section 7.5 tackled a topic of both theoretical and practical importance to the epistemological community: the use of online socio-technical systems (social machines) to generate collective knowledge. This is a topic that resonates with the interests of both social and applied epistemology [195, 380, 382].

These issues reveal the scope and scale of the research effort that will be required to advance our understanding of the epistemic implications of the Web. As with the discipline of Web Science, this research effort is very much an interdisciplinary endeavor, requiring collaboration between members of both the philosophical (e.g., epistemology) and scientific (e.g., computer science) communities. It is also a research effort that is poised to influence the design of future Web technologies. Inasmuch as we aim to press maximal epistemic benefit from the Web, both as the producers (see Section 7.5) and the consumers (see Section 4.5) of epistemic goods and services, it is essential that we provide a means for philosophical and scientific outputs to influence the design of Web-related technologies. Web Science is, of course, the natural home for this interdisciplinary endeavor. But perhaps we should go further than simply expand the existing intellectual remit of Web Science. Perhaps the relevant research effort is of sufficient scope and scale as to warrant

the creation of a new discipline. Such, at least, is the view of Palermos (Orestis Palermos, personal communication, January 28, 2016). He suggests that a new discipline—provisionally labeled *Web Epistemology*—should be established as a sub-discipline of Web Science. Hopefully, the present review goes some way to limning the general contours of this new region of the Web-related research landscape.

8.5 Ecological Engineering

Throughout this review, we have appealed to the idea that the Web serves as a form of ecosystem or ecology [460]—one that helps to shape cognitive processing at both an individual and collective level. Such a view establishes a natural point of contact with ecologically-oriented approaches in both cognitive science and the philosophy of mind [30, 256, 321, 512]. However, it is also a view that changes the way we think about human–Web interactions. From the standpoint of an ecological perspective, we have suggested that our interactions and engagements with the online environment should be seen as a form of ecological (or ecosystem) engineering. There are deep continuities here with the notion of *niche construction* in evolutionary theory [294, 366, 480, 482]. According to the proponents of niche construction, organisms actively engineer their local environments (niches) in ways that alter the kinds of selective pressure operating on future generations. This idea contributes to our understanding of the evolutionary mechanisms that may underlie the emergence of human cognitive capabilities [293, 480, 482]. But it is also possible to see ecological engineering as working to shape human cognitive capabilities on a more restricted timescale. Clark [102], for example, introduces us to the notion of *cognitive niche construction*, which he defines as “the process by which animals build physical structures that transform problem spaces in ways that aid (or sometimes impede) thinking and reasoning about some target domain or domains” (p. 62).³ When we combine these ideas with claims about

³The distinction between *cognitive niche construction* and *niche construction* is a somewhat artificial distinction, since some philosophers, most notably Kim Sterelny, also emphasize the cognitive relevance of niche construction processes. In the present context, the primary distinction between these terms relates to the period of time

the status of the Web as a form of cognitive ecology [see 460], the nature of our interactions and engagements with the Web take on a new significance. In particular, we can begin to see humanity as actively engaged in a cognitively-potent form of ecological engineering, progressively altering the specific mix of opportunities and affordances that shape the cognitive profiles of both ourselves and the generations that succeed us.

It is important to be clear what is meant by this particular proposal; for the reference to “ecological engineering” perhaps invites a somewhat limited view of the kinds of activity that are encompassed by the term. There is obviously a sense in which we can think of the activities of technology developers, software engineers and website designers as a form of ecological engineering (literally, engineering of the online environment). Similarly, it seems reasonable to think that the editing of a Wikipedia article or the uploading of a YouTube video should be regarded as forms of ecological engineering. But the notion of ecological engineering is actually broader than this. It includes, for example, the digital traces or ‘footprints’ that mark our treks through the online environment. In this sense, practically every form of interaction we have with the Web provides an opportunity for ecological engineering. Every time we perform a Google search, for example, we provide data that can be used for a variety of purposes (e.g., personalized search [439] or query-based syndromic surveillance [154]). The same applies to purchasing decisions on Amazon. Every time we make a purchase on Amazon, we provide information that can be used to influence the kind of informational contact that we (and others) have with the ‘Amazonian’ ecosystem [306]. The main point, here, is that the Web is a highly dynamic system, and every interaction we have with that system alters it in some way. As a result, we do not need to upload a video to YouTube in order to alter the online environment; for even the act of

over which they operate. We can thus distinguish between two kinds of epoch: the Clarkian (cognitive niche construction) and the Sterelnyian (niche construction). The Clarkian (based on the work of Andy Clark [102]) emphasizes the processes that operate in the here-and-now (i.e., on an *intra*-generational timescale). The Sterelnyian (based on the work of Kim Sterelny [480]), in contrast, emphasizes the processes that operate across generations (i.e., on an *inter*-generational timescale).

watching a YouTube video can alter the environment in such a way as to influence the cognitive responses of ourselves and others (recall the discussion of the socio-cognitive feedback hypothesis in Section 6.3). It is in this sense that we are all (potentially at least) ecological engineers of the online realm.

The cognitive significance of ecological engineering can be cashed out in a number of ways. It is perhaps most obviously thought of as a way of altering individual cognitive capabilities in the here-and-now, or at least as a way of altering cognition over relatively short (intra-generational) timescales. We encountered something along these lines in Section 3.5, where we discussed the relationship between self-tracking techniques and the technological transformation of the self. Such ideas dovetail with claims about the role of information and communication technologies in supporting the emergence of personal ecosystems [520] and informational ecologies for the self [176]. They also chime with claims about the role of ecological engineering in transforming our cognitive power and potential. As Clark [102] notes, “[i]n building our physical and social worlds, we build (or rather, we massively reconfigure) our minds and our capacities of thought and reason” (p. xxviii).

Another route to cognitive transformation follows a somewhat more circuitous trajectory. The main focus of attention, in this case, is the societal context in which human cognitive capabilities develop. A number of authors have drawn attention to the role of the Web and Internet in effecting various forms of social change [34, 324], and these ideas are further reinforced once we see the Web as a *social* technology, i.e., as a technology that lies at the heart of an ever-expanding array of social activities and social processes. Inasmuch as we see the Web as playing a crucial role in the computational realization of social processes [485]—as part of the material fabric that realizes social phenomena [464]—then it is easy to see how changes in the social environment could be causally and constitutively linked to changes in the online world. Ecological engineering (of the online environment) is, in this sense, a form of ‘social engineering’—a way of altering the structures, institutions, norms, cultural practices, and so on, that make our society what it is.

The cognitive significance of this claim comes into clearer view once we direct our attention to issues of niche construction and the socially-situated nature of human cognition. For the human social environment has often been implicated in the ontogenetic processes that give rise to human intelligence [see 305]. From this perspective, we might expect changes in society to alter the capabilities of future generations. To see this, we need only reflect on the way in which other kinds of technological development have helped to shape both society and ourselves. The invention of timekeeping devices, for example, led to profound forms of social, economic and intellectual change [297], with some suggesting that the transition from medieval clock towers to portable timekeeping devices played an important role in “propelling us out of the Middle Ages and into the Renaissance and then the Enlightenment” [81, p. 43]. Much the same, of course, could be said about the invention of orthographic practices and writing technology. The main point, here, is that social change often accompanies technological change, and social change leads to shifts in the forces and factors from which future minds emerge:

We engineer our own learning environments so as to create artificial developmental cocoons that impact our acquired capacities of thought and reason. Those enhanced minds then design new cognitive niches that train new generations of minds, and so on, in an empowering spiral of co-evolving complexity. [99, p. 278]

The concept of ecological engineering is thus an important one. When applied to the Web, we can see ecological engineering as not simply a means of changing the online environment (although it is surely that); we can also see it as a means of changing ourselves and the societies in which we live. From its humble beginnings as a tool to support the sharing of scientific data, the Web has emerged as a critical element of the broader cognitive ecology in which our biological brains are now situated. As we move into an era in which the Web plays an ever-more crucial role in shaping the course of our individual and collective endeavors, so we encounter new opportunities for cognitively-potent forms of interaction and engagement with the largest space of

knowledge and information our species has ever seen. It is, as yet, unclear exactly how the technological and social landscape of the Web will shape the contours of the human mind. However, as ecological engineers of the online world, we are at least in a position to influence the things that shape us. There can be little doubt that we are among the most capable of species when it comes to altering our environment—indeed, it seems that human activity has led to the creation of a new epoch of geological time [533, 560]. But in addition to being prodigious engineers, we are also among the most metamorphic of creatures. Just like the butterfly, our true nature is never really fixed. In changing our social and technological environments we shape the course of our cognitive destiny, altering our sense as to what we are, what we can do, and, perhaps most importantly, our sense as to what we might yet become.

9

Conclusion

We have now surveyed a wealth of literature that charts (at least in broad outline) the points of contact between Web Science, cognitive science and the philosophy of mind. There is, of course, much more that could have been done. We have not, for example, covered the considerable body of work relating to issues of cognitive development and cognitive ontogeny [38, 205, 206, 355, 356, 391]. Neither have we made any attempt to examine the Web from the standpoint of other theoretical positions within cognitive science, such as those afforded by grounded cognition [28, 29, 386], scaffolded cognition [483], and enactive cognition [502]. Another important omission relates to recent work on machine cognition and machine intelligence [452]. In this case, it has been suggested that the Web may have as much impact on machine cognition as it does on human cognition, with Web-based forms of contact with the human social environment helping to bring about state-of-the-art improvements in machine-based cognitive capabilities [451, 463].

In spite of these omissions, we hope the present review has provided at least some insight into the nature of the intellectual terrain that lies at the intersection of Web Science, cognitive science and the philosophy of mind. Understanding this terrain in more detail is an important goal

for Web Science. Given that the future of ourselves (and perhaps Planet Earth as a whole) is tied to the success or failure of our individual and collective cognitive engines, perhaps there are few other scientific undertakings that can claim to be as important.

Acknowledgements

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

Robert Clowes: This work is supported by the Portuguese Foundation for Science and Technology (FCT) under grant number SFRH/BPD/70440/2010 on “Virtualism and the Mind: Rethinking Presence, Representation and Self” and by an FCT strategic project grant for IFILNOVA, FCSH/NOVA UID/FIL/00183/2013.

All: The authors would like to thank two anonymous reviewers for their valuable insights and comments on an earlier version of the present review.

References

- [1] Adams, F. and K. Aizawa. 2001. “The bounds of cognition”. *Philosophical Psychology*. 14(1): 43–64.
- [2] Adams, F. and K. Aizawa. 2010. “Defending the bounds of cognition”. In: *The Extended Mind*. Ed. by R. Menary. Cambridge, Massachusetts, USA: MIT Press.
- [3] Adams, F. and R. Garrison. 2013. “The mark of the cognitive”. *Minds and Machines*. 23(3): 339–352.
- [4] Adams, F. 2010. “Why we still need a mark of the cognitive”. *Cognitive Systems Research*. 11(4): 324–331.
- [5] Adams, F. and K. Aizawa. 2008. *The Bounds of Cognition*. Oxford, UK: Blackwell.
- [6] Aharony, N., W. Pan, C. Ip, I. Khayal, and A. Pentland. 2011. “Social fMRI: Investigating and shaping social mechanisms in the real world”. *Pervasive and Mobile Computing*. 7(6): 643–659.
- [7] Aizawa, K. 2014. “Extended Cognition”. In: *The Routledge Handbook of Embodied Cognition*. Ed. by L. A. Shapiro. New York, New York, USA: Routledge.
- [8] Allen, C. in press. “On (not) defining cognition”. *Synthese*.
- [9] Allen, L. Q. 1995. “The effects of emblematic gestures on the development and access of mental representations of French expressions”. *The Modern Language Journal*. 79(4): 521–529.

- [10] Alter, A. L. and D. M. Oppenheimer. 2009. “Uniting the tribes of fluency to form a metacognitive nation”. *Personality and Social Psychology Review*. 13(3): 219–235.
- [11] Anderson, M. L. 2003. “Embodied cognition: A field guide”. *Artificial Intelligence*. 149(1): 91–130.
- [12] Andrighetto, G. and R. Conte. 2014. “Loops in Social Dynamics”. In: *Minding Norms: Mechanisms and Dynamics of Social Order in Agent Societies*. Ed. by R. Conte, G. Andrighetto, and M. Campenni. New York, New York, USA: Oxford University Press.
- [13] Aral, S. and D. Walker. 2012. “Identifying influential and susceptible members of social networks”. *Science*. 337(6092): 337–341.
- [14] Arango-Muñoz, S. 2014. “The nature of epistemic feelings”. *Philosophical Psychology*. 27(2): 193–211.
- [15] Araque, A. and M. Navarrete. 2010. “Glial cells in neuronal network function”. *Philosophical Transactions of the Royal Society B: Biological Sciences*. 365(1551): 2375–2381.
- [16] Arnau, E., S. Ayala, and T. Sturm. 2014. “Cognitive externalism meets bounded rationality”. *Philosophical Psychology*. 27(1): 50–64.
- [17] Auer, S., C. Bizer, G. Kobilarov, J. Lehmann, R. Cyganiak, and Z. Ives. 2007. “DBpedia: A nucleus for a web of open data”. *Lecture Notes in Computer Science*. 4825: 722–735.
- [18] Augustine, N. 1998. “What we don’t know does hurt us. How scientific illiteracy hobbles society”. *Science*. 279(5357): 1640–1641.
- [19] Bachrach, Y., M. Kosinski, T. Graepel, P. Kohli, and D. Stillwell. 2012. “Personality and patterns of Facebook usage”. In: *4th Annual ACM Web Science Conference*. Evanston, Illinois, USA.
- [20] Baddeley, A., M. W. Eysenck, and M. C. Anderson. 2009. *Memory*. Hove, East Sussex, UK: Psychology Press.
- [21] Baehr, J. 2012. *The Inquiring Mind: On Intellectual Virtues and Virtue Epistemology*. Oxford, UK: Oxford University Press.
- [22] Bainbridge, W. S. 2007. “The scientific research potential of virtual worlds”. *Science*. 317(5837): 472–476.

- [23] Banbury, S. and S. Tremblay, eds. 2004. *A Cognitive Approach to Situation Awareness: Theory and Application*. Aldershot, England, UK: Ashgate Publishing Company.
- [24] Barabási, A.-L. 2002. *Linked: The New Science of Networks*. Cambridge, Massachusetts, USA: Perseus Publishing.
- [25] Barabási, A.-L. 2016. *Network Science*. Cambridge, UK: Cambridge University Press.
- [26] Baronchelli, A., R. Ferrer-i-Cancho, R. Pastor-Satorras, N. Chater, and M. H. Christiansen. 2013. “Networks in Cognitive Science”. *Trends in Cognitive Sciences*. 17(7): 348–360.
- [27] Barrett, L. 2011. *Beyond the Brain: How Body and Environment Shape Animal and Human Minds*. Princeton, New Jersey, USA: Princeton University Press.
- [28] Barsalou, L. 2008. “Grounded cognition”. *Annual Review of Psychology*. 59: 617–645.
- [29] Barsalou, L. W. 2010. “Grounded Cognition: Past, Present, and Future”. *Topics in Cognitive Science*. 2(4): 716–724.
- [30] Bateson, G. 1972. *Steps to an Ecology of Mind*. London, UK: Jason Aronson Inc.
- [31] Battaly, H. 2008. “Virtue epistemology”. *Philosophy Compass*. 3(4): 639–663.
- [32] Bauckhage, C. and K. Kersting. 2016. “Collective attention on the Web”. *Foundation and Trends in Web Science*. 5(1–2): 1–136.
- [33] Bavelas, A. 1950. “Communication patterns in task-oriented groups”. *The Journal of the Acoustical Society of America*. 22(6): 725–730.
- [34] BBVA, ed. 2014. *Ch@nge: 19 Key Essays on How Internet Is Changing our Lives*. Madrid, Spain: Turner Publicaciones.
- [35] Bechtel, W., A. Abrahamsen, and G. Graham. 1998. “The life of cognitive science”. In: *A Companion to Cognitive Science*. Ed. by W. Bechtel and G. Graham. Malden, Massachusetts, USA: Blackwell Publishing Ltd.

- [36] Bechtel, W. and R. C. Richardson. 1993. *Discovering Complexity: Decomposition and Localization as Strategies in Scientific Research*. Princeton, New Jersey, USA: Princeton University Press.
- [37] Bell, G. and J. Gemmell. 2009. *Total Recall: How the E-Memory Revolution Will Change Everything*. London, UK: Dutton.
- [38] Bell, V., D. V. M. Bishop, and A. K. Przybylski. 2015. “The debate over digital technology and young people”. *The British Medical Journal*. 351: h3064.
- [39] Ben-Naim, J., J.-F. Bonnefon, A. Herzig, S. Leblois, and E. Lorini. 2013. “Computer-Mediated Trust in Self-interested Expert Recommendations”. In: *Cognition Beyond the Brain: Computation, Interactivity and Human Artifice*. Ed. by S. J. Cowley and F. Vallée-Tourangeau. London, UK: Springer-Verlag.
- [40] Berger, C. R., R. R. Gardner, M. R. Parks, L. Schulman, and G. R. Miller. 1976. “Interpersonal epistemology and interpersonal communication”. In: *Explorations in Interpersonal Communication*. Ed. by G. R. Miller. Beverly Hills, California, USA: Sage.
- [41] Berger, T. W. and D. L. Glanzman, eds. 2005. *Toward Replacement Parts for the Brain: Implantable Biomimetic Electronics as Neural Prostheses*. Cambridge, Massachusetts, USA: MIT Press.
- [42] Berger, T. W., R. E. Hampson, D. Song, A. Goonawardena, V. Z. Marmarelis, and S. A. Deadwyler. 2011. “A cortical neural prosthesis for restoring and enhancing memory”. *Journal of Neural Engineering*. 8(4): 046017.
- [43] Berners-Lee, T. and M. Fischetti. 1999. *Weaving the Web: The Original Design and Ultimate Destiny of the World Wide Web*. New York, New York, USA: Harper Collins.
- [44] Berners-Lee, T., W. Hall, J. A. Hendler, K. O’Hara, N. Shadbolt, and D. J. Weitzner. 2006a. “A framework for Web Science”. *Foundations and Trends in Web Science*. 1(1): 1–130.
- [45] Berners-Lee, T., W. Hall, J. Hendler, N. Shadbolt, and D. J. Weitzner. 2006b. “Creating a Science of the Web”. *Science*. 313(5788): 769.

- [46] Berners-Lee, T., J. Hendler, and O. Lassila. 2001. “The Semantic Web”. *Scientific American*. 284(4): 34–43.
- [47] Bernstein, A., M. Klein, and T. W. Malone. 2012. “Programming the global brain”. *Communications of the ACM*. 55(5): 41–43.
- [48] Bernstein, M. S., G. Little, R. C. Miller, B. Hartmann, M. S. Ackerman, D. R. Karger, D. Crowell, and K. Panovich. 2015. “Soylent: A word processor with a crowd inside”. *Communications of the ACM*. 58(8): 85–94.
- [49] Berry, E., N. Kapur, L. Williams, S. Hodges, P. Watson, G. Smyth, J. Srinivasan, R. Smith, B. Wilson, and K. Wood. 2007. “The use of a wearable camera, SenseCam, as a pictorial diary to improve autobiographical memory in a patient with limbic encephalitis: A preliminary report”. *Neuropsychological Rehabilitation*. 17(4): 582–601.
- [50] Bigham, J. P., C. Jayant, H. Ji, G. Little, A. Miller, R. C. Miller, R. Miller, A. Tatarowicz, B. White, and S. White. 2010. “VizWiz: Nearly real-time answers to visual questions”. In: *23rd Annual ACM Symposium on User Interface Software and Technology*. New York, New York, USA.
- [51] Bikakis, N. and T. Sellis. 2016. “Exploration and visualization in the web of big linked data: A survey of the state of the art”. In: *16th International Workshop on Linked Data Management*. Bordeaux, France.
- [52] Biocca, F. A. and J. P. Rolland. 1998. “Virtual eyes can rearrange your body: Adaptation to visual displacement in see-through, head-mounted displays”. *Presence: Teleoperators and Virtual Environments*. 7(3): 262–277.
- [53] Bizer, C. 2009. “The emerging web of linked data”. *IEEE Intelligent Systems*. 24(5): 87–92.
- [54] Bizer, C., T. Heath, and T. Berners-Lee. 2009a. “Linked data—the story so far”. *International Journal on Semantic Web and Information Systems*. 5(3): 1–22.

- [55] Bizer, C., J. Lehmann, G. Kobilarov, S. Auer, C. Becker, R. Cyganiak, and S. Hellmann. 2009b. “DBpedia—A crystallization point for the Web of Data”. *Web Semantics: Science, Services and Agents on the World Wide Web*. 7(3): 154–165.
- [56] Bjerring, J. C. and N. J. L. L. Pedersen. 2014. “All the (many, many) things we know: Extended knowledge”. *Philosophical Issues*. 24(1): 24–38.
- [57] Blomberg, O. 2011. “Conceptions of Cognition for Cognitive Engineering”. *The International Journal of Aviation Psychology*. 21(1): 85–104.
- [58] Bode, M. and D. B. Kristensen. 2016. “The digital doppelgänger within. A study on self-tracking and the quantified self movement”. In: *Assembling Consumption: Researching Actors, Networks and Markets*. Ed. by R. Canniford and D. Bajde. New York, New York, USA: Routledge.
- [59] Bonabeau, E. 2009. “Decisions 2.0: The Power of Collective Intelligence”. *MIT Sloan Management Review*. 50(2): 45–52.
- [60] Bonabeau, E., M. Dorigo, and G. Theraulaz. 1999. *Swarm Intelligence: From Natural to Artificial Systems*. New York, New York, USA: Oxford University Press.
- [61] Borgatti, S. P., A. Mehra, D. J. Brass, and G. Labianca. 2009. “Network analysis in the social sciences”. *Science*. 323(5916): 892–895.
- [62] Bott, E. 2001. *Family and Social Network: Roles, Norms and External Relationships in Ordinary Urban Families*. 2nd. Abingdon, Oxon, UK: Routledge.
- [63] Boulianne, S. 2009. “Does Internet use affect engagement? A meta-analysis of research”. *Political communication*. 26(2): 193–211.
- [64] Braitenberg, V. 1984. *Vehicles: Experiments in Synthetic Psychology*. Cambridge, Massachusetts, USA: MIT Press.
- [65] Briggles, A. 2008. “Real friends: How the Internet can foster friendship”. *Ethics and Information Technology*. 10(1): 71–79.

- [66] Brin, S. and L. Page. 1998. “The anatomy of a large-scale hyper-textual search engine”. In: *7th Annual Conference of the World Wide Web*. Brisbane, Australia.
- [67] Bruggeman, H., C. Kliman-Silver, F. Domini, and J.-H. Song. 2013. “Dynamic manipulation generates touch information that can modify vision”. *Psychological Science*. 24(6): 1063–1065.
- [68] Buchanan, M. 2002. *Nexus: Small Worlds and the Groundbreaking Science of Networks*. New York, New York, USA: W. W. Norton & Company.
- [69] Buck, R. 1980. “Nonverbal behavior and the theory of emotion: The facial feedback hypothesis”. *Journal of Personality and Social Psychology*. 38(5): 811–824.
- [70] Burégio, V. A., S. R. Meira, N. S. Rosa, and V. C. Garcia. 2013. “Moving towards “relationship-aware” applications and services: A social machine-oriented approach”. In: *6th International Workshop on Evolutionary Business Processes*. Vancouver, British Columbia, Canada.
- [71] Burkell, J. A. 2016. “Remembering me: Big Data, individual identity, and the psychological necessity of forgetting”. *Ethics and Information Technology*. 18(1): 17–23.
- [72] Byrne, R. W. and A. Whiten. 1988. *Machiavellian Intelligence: Social Expertise and the Evolution of Intellect in Monkeys, Apes, and Humans*. Oxford, UK: Oxford University Press.
- [73] Calvo, P. and F. Keijzer. 2009. “Cognition in plants”. In: *Plant-Environment Interactions: From Sensory Plant Biology to Active Plant Behavior*. Ed. by F. Baluska. Berlin, Germany: Springer-Verlag.
- [74] Campbell, J. I. D., ed. 2005. *Handbook of Mathematical Cognition*. Hove, East Sussex, UK: Psychology Press.
- [75] Canny, J. and E. Paulos. 2000. “Tele-embodiment and shattered presence: Reconstructing the body for online interaction”. In: *The Robot in the Garden: Telerobotics and Telepistemology in the Age of the Internet*. Ed. by K. Goldberg. Cambridge, Massachusetts, USA: MIT Press.

- [76] Cardamone, C., K. Schawinski, M. Sarzi, S. P. Bamford, N. Bennert, C. M. Urry, C. Lintott, W. C. Keel, J. Parejko, and R. C. Nichol. 2009. “Galaxy Zoo Green Peas: Discovery of a class of compact extremely star-forming galaxies”. *Monthly Notices of the Royal Astronomical Society*. 399(3): 1191–1205.
- [77] Cardinali, L., F. Frassinetti, C. Brozzoli, C. Urquizar, A. C. Roy, and A. Farnè. 2009. “Tool-use induces morphological updating of the body schema”. *Current Biology*. 19(12): R478–R479.
- [78] Cardinali, L., S. Jacobs, C. Brozzoli, F. Frassinetti, A. C. Roy, and A. Farnè. 2012. “Grab an object with a tool and change your body: Tool-use-dependent changes of body representation for action”. *Experimental Brain Research*. 218(2): 259–271.
- [79] Carlson, T. A., G. Alvarez, D.-a. Wu, and F. A. J. Verstraten. 2010. “Rapid assimilation of external objects into the body schema”. *Psychological Science*. 21(7): 1000–1005.
- [80] Carr, N. 2008. “Is Google Making Us Stupid?” *The Atlantic*. 302(1): 56–63.
- [81] Carr, N. 2010. *The Shallows: How the Internet Is Changing the Way We Think, Read and Remember*. London, UK: Atlantic Books.
- [82] Carr, N. 2014. *The Glass Cage: How our Computers Are Changing Us*. New York, New York, USA: W. W. Norton & Company.
- [83] Carter, A. J. and S. O. Palermos. 2016. “The ethics of extended cognition: Is having your computer compromised a personal assault?” *Journal of the American Philosophical Association*. 2(4): 542–560.
- [84] Castelfranchi, C., M. Piunti, A. Ricci, and L. Tummolini. 2012. “AmI Systems as Agent-Based Mirror Worlds: Bridging Humans and Agents through Stigmergy”. In: *Agents and Ambient Intelligence*. Ed. by T. Bosse. Amsterdam, The Netherlands: IOS Press.
- [85] Castells, M. 2010. *The Rise of the Network Society*. 2nd. Chichester, UK: Wiley-Blackwell.
- [86] Cate, C. ten and S. D. Healy, eds. 2017. *Avian Cognition*. Cambridge, UK: Cambridge University Press.

- [87] Caton, S., M. Hall, and C. Weinhardt. 2015. "How do politicians use Facebook? An applied Social Observatory". *Big Data & Society*. 2(2).
- [88] Chase, J. and D. Coady. in press. *The Routledge Handbook of Applied Epistemology*. New York, New York, USA: Routledge.
- [89] Cheney-Lippold, J. 2011. "A new algorithmic identity: Soft biopolitics and the modulation of control". *Theory, Culture & Society*. 28(6): 164–181.
- [90] Chi, E. H. 2008. "The Social Web: Research and opportunities". *Computer*. 41(9): 88–91.
- [91] Chi, E. H. 2009. "Augmented social cognition: Using Social Web technology to enhance the ability of groups to remember, think, and reason". In: *35th SIGMOD International Conference on Management of Data*. Providence, Rhode Island, USA.
- [92] Chi, E. H., P. Pirolli, B. Suh, A. Kittur, B. Pendleton, and T. Mytkowicz. 2008. "Augmented Social Cognition". In: *AAAI Spring Symposium on Social Information Processing*. Stanford, California, USA.
- [93] Chrisley, R. 2003. "Embodied artificial intelligence". *Artificial Intelligence*. 149(1): 131–150.
- [94] Christakis, N. A. and J. H. Fowler. 2007. "The spread of obesity in a large social network over 32 years". *New England Journal of Medicine*. 357(4): 370–379.
- [95] Christakis, N. and J. Fowler. 2009. *Connected: The Surprising Power of Social Networks and How They Shape Our Lives*. London, UK: Little, Brown and Company.
- [96] Clark, A. 1997. *Being There: Putting Brain, Body and World Together Again*. Cambridge, Massachusetts, USA: MIT Press.
- [97] Clark, A. 1999. "Where brain, body, and world collide". *Cognitive Systems Research*. 1(1): 5–17.
- [98] Clark, A. 2003. *Natural-Born Cyborgs: Minds, Technologies and the Future of Human Intelligence*. Oxford, UK: Oxford University Press.

- [99] Clark, A. 2007a. “Re-Inventing Ourselves: The Plasticity of Embodiment, Sensing, and Mind”. *Journal of Medicine and Philosophy*. 32(3): 263–282.
- [100] Clark, A. 2007b. “Soft Selves and Ecological Control”. In: *Distributed Cognition and the Will: Individual Volition and Social Context*. Ed. by D. Ross, D. Spurrett, H. Kincaid, and G. L. Stephens. Cambridge, Massachusetts, USA: MIT Press.
- [101] Clark, A. 2008a. “Pressing the Flesh: A Tension in the Study of the Embodied, Embedded Mind?” *Philosophy and Phenomenological Research*. 76(1): 37–59.
- [102] Clark, A. 2008b. *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. New York, New York, USA: Oxford University Press.
- [103] Clark, A. 2009. “Dispersed Selves”. *Leonardo Electronic Almanac*. 16(4–5).
- [104] Clark, A. 2010a. “Coupling, Constitution and the Cognitive Kind: A Reply to Adams and Aizawa”. In: *The Extended Mind*. Ed. by R. Menary. Cambridge, Massachusetts, USA: MIT Press.
- [105] Clark, A. 2010b. “Memento’s revenge: The extended mind, extended”. In: *The Extended Mind*. Ed. by R. Menary. Cambridge, Massachusetts, USA: MIT Press.
- [106] Clark, A. 2011. “Finding the mind”. *Philosophical Studies*. 152(3): 447–461.
- [107] Clark, A. 2015. “What ‘Extended Me’ knows”. *Synthese*. 192(11): 3757–3775.
- [108] Clark, A. and D. Chalmers. 1998. “The Extended Mind”. *Analysis*. 58(1): 7–19.
- [109] Clowes, R. W. 2012. “Hybrid Memory, Cognitive Technology and Self”. In: *Proceedings of the AISB (The Society for the Study of Artificial Intelligence and Simulation of Behaviour) and IACAP (The International Association for Computing and Philosophy) World Congress*. Birmingham, UK.
- [110] Clowes, R. W. 2013. “The Cognitive Integration of E-Memory”. *Review of Philosophy and Psychology*. 4(1): 107–133.

- [111] Clowes, R. W. 2014a. “Faceache: WEB 2.0, Safety Culture, The End of Intimacy and the Implosion of Private Life”. In: *Público/Privado. O Deslizar de uma Fronteira*. Ed. by F. Negro. Lisbon, Portugal: Húmus/CEC.
- [112] Clowes, R. W. 2014b. “Reading, New Media Technology and what they tell us about Social Cognition”. In: *European Conference on Social Intelligence*. Barcelona, Spain.
- [113] Clowes, R. W. 2015. “Thinking in the Cloud: The Cognitive Incorporation of Cloud-Based Technology”. *Philosophy & Technology*. 28(2): 261–296.
- [114] Clowes, R. W. in press(a). “Extended Memory”. In: *The Routledge Handbook of Philosophy of Memory*. Ed. by S. Bernecker and K. Michaelian. New York, New York, USA: Routledge.
- [115] Clowes, R. W. in press(b). “Screen Reading and the Creation of New Cognitive Ecologies”. *AI & Society*.
- [116] Clowes, R. W. Submitted. “Immaterial Engagement: The New Ecology of Human Agency and the Artefactual Turn in the Cognitive Sciences”.
- [117] Coady, D. 2012. *What to Believe Now: Applying Epistemology to Contemporary Issues*. Oxford, UK: Blackwell.
- [118] Coburn, C. 2014. “Play to Cure: Genes in Space”. *Lancet Oncology*. 15(7): 688.
- [119] Cohen, J. 2013. “Memory Implants”. *MIT Technology Review*. 116(3): 62–68.
- [120] Colombo, M. 2017. “Why build a virtual brain? Large-scale neural simulations as jump start for cognitive computing”. *Journal of Experimental & Theoretical Artificial Intelligence*. 29(2): 361–370.
- [121] Committee on Network Science for Future Army Applications. 2005. *Network Science*. Washington D.C., USA: The National Academies Press.

- [122] Conte, R., G. Andrighetto, M. Campennì, and M. Paolucci. 2007. “Emergent and immergent effects in complex social systems”. In: *AAAI Fall Symposium: Emergent Agents and Socialities: Social and Organizational Aspects of Intelligence*. Arlington, Virginia, USA.
- [123] Conte, R., N. Gilbert, G. Bonelli, C. Cioffi-Revilla, G. Deffuant, J. Kertesz, V. Loreto, S. Moat, J.-P. Nadal, S. Sanchez, A. Nowak, A. Flache, M. San Miguel, and D. Helbing. 2012. “Manifesto of computational social science”. *The European Physical Journal Special Topics*. 214(1): 325–346.
- [124] Convertino, G., D. Billman, P. Pirolli, J. P. Massar, and J. Shrager. 2008. “The CACHE study: Group effects in computer-supported collaborative analysis”. *Computer Supported Cooperative Work*. 17(4): 353–393.
- [125] Cook, S. W., Z. Mitchell, and S. Goldin-Meadow. 2008. “Gesturing makes learning last”. *Cognition*. 106(2): 1047–1058.
- [126] Cook, S. W., T. K. Yip, and S. Goldin-Meadow. 2010. “Gesturing makes memories that last”. *Journal of Memory and Language*. 63(4): 465–475.
- [127] Cook, S. W., T. K. Yip, and S. Goldin-Meadow. 2012. “Gestures, but not meaningless movements, lighten working memory load when explaining math”. *Language and Cognitive Processes*. 27(4): 594–610.
- [128] Cooke, N. J., J. C. Gorman, and L. J. Rowe. 2004. “An ecological perspective on team cognition”. In: *Team Effectiveness in Complex Organizations: Cross-Disciplinary Perspectives and Approaches*. Ed. by E. Salas, G. F. Goodwin, and C. S. Burke. New York, New York, USA: Psychology Press.
- [129] Cooke, N. J., J. C. Gorman, and J. L. Winner. 2007. “Team Cognition”. In: *Handbook of Applied Cognition*. Ed. by F. T. Durso, R. S. Nickerson, S. T. Dumais, S. Lewandowsky, and T. J. Perfect. 2nd. Chichester, UK: John Wiley & Sons.

- [130] Cooper, S., F. Khatib, A. Treuille, J. Barbero, J. Lee, M. Beenen, A. Leaver-Fay, D. Baker, Z. Popović, and Foldit Players. 2010. “Predicting protein structures with a multiplayer online game”. *Nature*. 466(7307): 756–760.
- [131] Couzin, I. D. 2009. “Collective cognition in animal groups”. *Trends in Cognitive Sciences*. 13(1): 36–43.
- [132] Craver, C. 2007. “Constitutive explanatory relevance”. *Journal of Philosophical Research*. 32: 3–20.
- [133] Craver, C. F. and L. Darden. 2013. *In Search of Mechanisms: Discoveries Across the Life Sciences*. Chicago, Illinois, USA: The University of Chicago Press.
- [134] Craver, C. and J. Tabery. 2016. “Mechanisms in science”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by E. N. Zalta. Spring 2016. Stanford, California, USA: Stanford University.
- [135] Crouser, R. J., A. Ottley, and R. Chang. 2013. “Balancing Human and Machine Contributions in Human Computation Systems”. In: *Handbook of Human Computation*. Ed. by P. Michelucci. New York, New York, USA: Springer.
- [136] Dangerfield, J. M., T. S. McCarthy, and W. N. Ellery. 1998. “The mound-building termite *Macrotermes michaelseni* as an ecosystem engineer”. *Journal of Tropical Ecology*. 14(4): 507–520.
- [137] Darmailacq, A.-S., D. Ludovic, and J. Mather, eds. 2014. *Cephalopod Cognition*. Cambridge, UK: Cambridge University Press.
- [138] Dehaene, S. and L. Cohen. 2007. “Cultural recycling of cortical maps”. *Neuron*. 56(2): 384–398.
- [139] Demartini, G. 2015. “Hybrid human–machine information systems: Challenges and opportunities”. *Computer Networks*. 90: 5–13.
- [140] DeVaul, R. W., A. Pentland, and V. R. Corey. 2005. “The memory glasses: Subliminal vs. overt memory support with imperfect information”. In: *7th IEEE International Symposium on Wearable Computers*. White Plains, New York, USA.
- [141] Di Paolo, E. and E. Thompson. 2014. “The Enactive Approach”. In: *The Routledge Handbook of Embodied Cognition*. Ed. by L. A. Shapiro. New York, New York, USA: Routledge.

- [142] Diehl, M. and W. Stroebe. 1987. "Productivity loss in brainstorming groups: Toward the solution of a riddle". *Journal of Personality and Social Psychology*. 53: 497–509.
- [143] Dijk, J. van. 2007. *Mediated Memories in the Digital Age*. Stanford, California, USA: Stanford University Press.
- [144] Dijk, J. A. G. M. van. 2012. *The Network Society*. 3rd. London, UK: Sage Publications Ltd.
- [145] Dijk, J. A. G. M. van and A. J. A. M. van Deursen. 2014. *Digital Skills: Unlocking the Information Society*. New York, New York, USA: Palgrave Macmillan.
- [146] Dijkstra, K., M. P. Kaschak, and R. A. Zwaan. 2007. "Body posture facilitates retrieval of autobiographical memories". *Cognition*. 102(1): 139–149.
- [147] Dolins, F. L. and R. W. Mitchell, eds. 2010. *Spatial Cognition, Spatial Perception: Mapping the Self and Space*. Cambridge, UK: Cambridge University Press.
- [148] Donald, M. 1991. *Origins of the Modern Mind: Three Stages in the Evolution of Culture and Cognition*. Cambridge, Massachusetts, USA: Harvard University Press.
- [149] Donath, Judith and boyd, danah. 2004. "Public displays of connection". *BT Technology Journal*. 22(4): 71–82.
- [150] Dourish, P. 2001. *Where the Action Is: The Foundations of Embodied Interaction*. Cambridge, Massachusetts, USA: MIT Press.
- [151] Drain, C. and R. C. Strong. 2015. "Situated Mediation and Technological Reflexivity: Smartphones, Extended Memory, and Limits of Cognitive Enhancement". In: *Social Epistemology and Technology: Toward Public Self-Awareness Regarding Technological Mediation*. Ed. by F. Scalabrino. Lanham, Maryland, USA: Rowman & Littlefield International.
- [152] Dror, I. E. and S. Harnad. 2008. "Offloading cognition onto cognitive technology". In: *Cognition Distributed: How Cognitive Technology Extends Our Minds*. Ed. by I. E. Dror and S. Harnad. Amsterdam, The Netherlands: John Benjamins.

- [153] Dubourg, L., A. R. Silva, C. Fitamen, C. J. A. Moulin, and C. Souchay. 2016. "SenseCam: A new tool for memory rehabilitation?" *Revue Neurologique*. 172(12): 735–747.
- [154] Dugas, A. F., M. Jalalpour, Y. Gel, S. Levin, F. Torcaso, T. Igusa, and R. E. Rothman. 2013. "Influenza forecasting with Google Flu Trends". *PLoS One*. 8(2): e56176.
- [155] Dunbar, R. 2007. "Evolution of the social brain". In: *The Evolution of Mind: Fundamental Questions and Controversies*. Ed. by S. W. Gangestad and J. A. Simpson. New York, New York, USA: The Guilford Press.
- [156] Dunbar, R. I. M. 1992. "Neocortex size as a constraint on group size in primates". *Journal of Human Evolution*. 22(6): 469–493.
- [157] Dunbar, R. I. M. 1993. "Coevolution of neocortical size, group size and language in humans". *Behavioral and Brain Sciences*. 16(4): 681–735.
- [158] Dunbar, R. I. M. 1998. "The Social Brain Hypothesis". *Evolutionary Anthropology*. 6(5): 178–190.
- [159] Dunbar, R. I. M. 2012. "Social cognition on the Internet: Testing constraints on social network size". *Philosophical Transactions of the Royal Society B: Biological Sciences*. 367(1599): 2192–2201.
- [160] Dunbar, R. I. M. 2016. "Do online social media cut through the constraints that limit the size of offline social networks?" *Royal Society Open Science*. 3(1): 150292.
- [161] Dunn, T. L. and E. F. Risko. 2016. "Toward a Metacognitive Account of Cognitive Offloading". *Cognitive Science*. 40(5): 1080–1127.
- [162] Dyson, M. P., L. Hartling, J. Shulhan, A. Chisholm, A. Milne, P. Sundar, S. D. Scott, and A. S. Newton. 2016. "A Systematic Review of Social Media Use to Discuss and View Deliberate Self-Harm Acts". *PLoS One*. 11(5): e0155813.
- [163] Eck, D. van and H. L. de Jong. 2016. "Mechanistic explanation, cognitive systems demarcation, and extended cognition". *Studies in History and Philosophy of Science Part A*. 59: 11–21.
- [164] Endsley, M. R. 1995. "Toward a theory of situation awareness in dynamic systems". *Human factors*. 37(1): 32–64.

- [165] Eysenck, M. W. and M. T. Keane. 2010. *Cognitive Psychology: A Student's Handbook*. 6th. Hove, East Sussex, UK: Psychology Press.
- [166] Fallis, D. 2008. "Toward an epistemology of Wikipedia". *Journal of the American Society for Information Science and Technology*. 59(10): 1662–1674.
- [167] Fallis, D. 2011. "Wikipistemology". In: *Social Epistemology: Essential Readings*. Ed. by A. I. Goldman and D. Whitcomb. New York, New York, USA: Oxford University Press.
- [168] Fan, X. and J. Yen. 2007. "R-CAST: Integrating team intelligence for human-centered teamwork". In: *22nd Conference on Artificial Intelligence (AAAI-07)*. Vancouver, British Columbia, Canada.
- [169] Ferguson, A. M., D. McLean, and E. F. Risko. 2015. "Answers at your fingertips: Access to the Internet influences willingness to answer questions". *Consciousness and Cognition*. 37: 91–102.
- [170] Ferrara, E., O. Varol, C. Davis, F. Menczer, and A. Flammini. 2016. "The rise of social bots". *Communications of the ACM*. 59(7): 96–104.
- [171] Ferrucci, D., E. Brown, J. Chu-Carroll, J. Fan, D. Gondek, A. A. Kalyanpur, A. Lally, J. W. Murdock, E. Nyberg, and J. Prager. 2010. "Building Watson: An overview of the DeepQA project". *AI Magazine*. 31(3): 59–79.
- [172] Finzi, E. 2013. *The Face of Emotion: How Botox Affects Our Moods and Relationships*. New York, New York, USA: Palgrave Macmillan.
- [173] Fisher, M., M. K. Goddu, and F. C. Keil. 2015. "Searching for Explanations: How the Internet Inflates Estimates of Internal Knowledge". *Journal of Experimental Psychology: General*. 144(3): 674–687.
- [174] Fitzsimmons, G., M. Weal, and D. Drieghe. 2014. "Skim reading: An adaptive strategy for reading on the Web". In: *6th ACM Conference on Web Science*. Bloomington, Indiana, USA.
- [175] Floridi, L. 2011a. "The Construction of Personal Identities Online". *Minds and Machines*. 21(4): 477–479.

- [176] Floridi, L. 2011b. “The Informational Nature of Personal Identity”. *Minds and Machines*. 21(4): 549–566.
- [177] Floridi, L. 2014. *The 4th Revolution: How the Infosphere is Reshaping Human Reality*. Oxford, UK: Oxford University Press.
- [178] Fogg, B. J., J. Marshall, O. Laraki, A. Osipovich, C. Varma, N. Fang, J. Paul, A. Rangnekar, J. Shon, P. Swani, and M. Treinen. 2001. “What makes Web sites credible?: A report on a large quantitative study”. In: *Conference on Human Factors in Computing Systems*. Seattle, Washington.
- [179] Fogg, B. J., C. Soohoo, D. R. Danielson, L. Marable, J. Stanford, and E. R. Tauber. 2003. “How do users evaluate the credibility of Web sites? A study with over 2,500 participants”. In: *Conference on Designing for User Experiences*. San Francisco, California, USA.
- [180] Foroughi, J., T. Mitew, P. Ogunbona, R. Raad, and F. Safaei. 2016. “Smart Fabrics and networked clothing: Recent developments in CNT-based fibers and their continual refinement”. *IEEE Consumer Electronics Magazine*. 5(4): 105–111.
- [181] Forsyth, D. R. 2006. *Group Dynamics*. Thomas Wadsworth: Belmont, California, USA.
- [182] Gallagher, S. 2013. “The socially extended mind”. *Cognitive Systems Research*. 25–26: 4–12.
- [183] Gant, D. and S. Kiesler. 2001. “Blurring the boundaries: Cell phones, mobility, and the line between work and personal life”. In: *Wireless World: Social and Interactional Aspects of the Mobile Age*. Ed. by B. Brown, N. Green, and R. Harper. London, UK: Springer.
- [184] Geiger, R. S. and D. Ribes. 2010. “The work of sustaining order in Wikipedia: The banning of a vandal”. In: *ACM Conference on Computer Supported Cooperative Work*. Savannah, Georgia, USA.
- [185] Giannotti, F., D. Pedreschi, A. Pentland, P. Lukowicz, D. Kossmann, J. Crowley, and D. Helbing. 2012. “A planetary nervous system for social mining and collective awareness”. *The European Physical Journal Special Topics*. 214(1): 49–75.

- [186] Gibbs, R. W. 2005. *Embodiment and Cognitive Science*. New York, New York, USA: Cambridge University Press.
- [187] Goga, O., G. Venkatadri, and K. P. Gummadi. 2015. “The doppelgänger bot attack: Exploring identity impersonation in online social networks”. In: *ACM Internet Measurement Conference*. Tokyo, Japan.
- [188] Golbeck, J. 2006. “Trust on the World Wide Web: A survey”. *Foundations and Trends in Web Science*. 1(2): 131–197.
- [189] Golbeck, J. 2013. *Analyzing the Social Web*. Waltham, Massachusetts, USA: Morgan Kaufmann.
- [190] Golbeck, J., C. Robles, M. Edmondson, and K. Turner. 2011. “Predicting personality from Twitter”. In: *IEEE Third International Conference on Social Computing (SocialCom)*. Boston, Massachusetts, USA.
- [191] Goldberg, S. 2011. “Self-Knowledge”. In: *The Routledge Companion to Epistemology*. Ed. by S. Bernecker and D. Pritchard. New York, New York, USA: Routledge.
- [192] Goldin-Meadow, S. 2003. *Hearing Gesture: How Our Hands Help Us Think*. Cambridge, Massachusetts, USA: Harvard University Press.
- [193] Goldin-Meadow, S. and S. L. Beilock. 2010. “Action’s influence on thought: The case of gesture”. *Perspectives on Psychological Science*. 5(6): 664–674.
- [194] Goldman, A. I. 2008. “The Social Epistemology of Blogging”. In: *Information Technology and Moral Philosophy*. Ed. by J. van den Hoven and J. Weckert. New York, New York, USA: Cambridge University Press.
- [195] Goldman, A. I. 2011. “A Guide to Social Epistemology”. In: *Social Epistemology: Essential Readings*. Ed. by A. I. Goldman and D. Whitcomb. New York, New York, USA: Oxford University Press.
- [196] Goldman, A. I. 2012. *Reliabilism and Contemporary Epistemology: Essays*. New York, New York, USA: Oxford University Press.

- [197] Goldstone, R. L. and G. Lupyan. 2016. “Discovering psychological principles by mining naturally occurring data sets”. *Topics in Cognitive Science*. 8(3): 548–568.
- [198] Gollakota, S., M. S. Reynolds, J. R. Smith, and D. J. Wetherall. 2014. “The emergence of RF-powered computing”. *Computer*. 47(1): 32–39.
- [199] González-Bailón, S., J. Borge-Holthoefery, and Y. Morenoy. 2014. “Online networks and the diffusion of protest”. In: *Analytical Sociology: Actions and Networks*. Ed. by G. Manzo. Chichester, West Sussex, UK: John Wiley & Sons, Ltd.
- [200] Good, B. M. and A. I. Su. 2011. “Games with a scientific purpose”. *Genome Biology*. 12(135): 1–3.
- [201] Gourine, A. V., V. Kasymov, N. Marina, F. Tang, M. F. Figueiredo, S. Lane, A. G. Teschemacher, K. M. Spyer, K. Deisseroth, and S. Kasparov. 2010. “Astrocytes control breathing through pH-dependent release of ATP”. *Science*. 329(5991): 571–575.
- [202] Goyal, S. 2014. “Public vs private vs hybrid vs community-cloud computing: A critical review”. *International Journal of Computer Network and Information Security*. 6(3): 20–29.
- [203] Greco, J. 2010. *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity*. Cambridge, UK: Cambridge University Press.
- [204] Greco, J. and J. Turri, eds. 2012. *Virtue Epistemology: Contemporary Readings*. Cambridge, Massachusetts, USA: MIT Press.
- [205] Greenfield, S. 2003. *Tomorrow’s People: How 21st Century Technology is Changing the Way We Think and Feel*. London, England: Penguin Books Ltd.
- [206] Greenfield, S. 2014. *Mind Change: How Digital Technologies Are Leaving Their Mark on Our Brains*. London, England: Rider.
- [207] Greengard, S. 2015. *The Internet of Things*. Cambridge, Massachusetts, USA: MIT Press.
- [208] Gregory, R. L. 1993. *Mind in Science: A History of Explanations in Psychology and Physics*. London, England: Penguin Books Ltd.

- [209] Groen, J. and P. J. Werkhoven. 1998. “Visuomotor adaptation to virtual hand position in interactive virtual environments”. *Presence: Teleoperators and Virtual Environments*. 7(5): 429–446.
- [210] Guinard, D. D. and V. M. Trifa. 2016. *Building the Web of Things*. Shelter Island, New York, USA: Manning Publications Co.
- [211] Guinard, D. and V. Trifa. 2009. “Towards the Web of Things: Web mashups for embedded devices”. In: *18th International World Wide Web Conference*. Madrid, Spain.
- [212] Guinard, D., V. Trifa, F. Mattern, and E. Wilde. 2011. “From the Internet of Things to the Web of Things: Resource-oriented Architecture and Best Practices”. In: *Architecting the Internet of Things*. Ed. by D. Uckelmann, M. Harrison, and F. Michahelles. Berlin, Germany: Springer-Verlag.
- [213] Gupta, A., H. Lamba, and P. Kumaraguru. 2013. “\$1.00 per RT #BostonMarathon #PrayForBoston: Analyzing fake content on Twitter”. In: *eCrime Researchers Summit*. San Francisco, California, USA.
- [214] Halpin, H. 2013. “Does the Web extend the mind?” In: *5th Annual ACM Web Science Conference*. Paris, France.
- [215] Halpin, H., A. Clark, and M. Wheeler. 2010. “Towards a Philosophy of the Web: Representation, Enaction, Collective Intelligence”. In: *Web Science Conference*. Raleigh, North Carolina, USA.
- [216] Halpin, H. and A. Monnin, eds. 2014. *Philosophical Engineering: Toward a Philosophy of the Web*. Chichester, West Sussex, UK: John Wiley & Sons Ltd.
- [217] Halpin, H. and A. Monnin. 2016. “The decentralization of knowledge: How Carnap and Heidegger influenced the Web”. *First Monday*. 21(12).
- [218] Hamilton, D. L., ed. 2005. *Social Cognition*. New York, New York, USA: Psychology Press.

- [219] Han, J. and M. Orshansky. 2013. “Approximate computing: An emerging paradigm for energy-efficient design”. In: *18th IEEE European Test Symposium*. Avignon, France.
- [220] Harnad, S. 1999. “The future of scholarly skywriting”. In: *I in the Sky: Visions of the Information Future*. Ed. by A. Scammell. London, UK: Aslib.
- [221] Harnad, S. and I. E. Dror. 2006. “Distributed cognition: Cognizing, autonomy and the Turing Test”. *Pragmatics & Cognition*. 14(2): 209–213.
- [222] Harvey, M., M. Langheinrich, and G. Ward. 2016. “Remembering through lifelogging: A survey of human memory augmentation”. *Pervasive and Mobile Computing*. 27: 14–26.
- [223] Hatfield, G. 2014. “Cognition”. In: *The Routledge Handbook of Embodied Cognition*. Ed. by L. A. Shapiro. New York, New York, USA: Routledge.
- [224] Hattori, S., T. Tezuka, and K. Tanaka. 2006. “Activity-based query refinement for context-aware information retrieval”. In: *9th International Conference on Asian Digital Libraries*. Kyoto, Japan.
- [225] Heath, J. and J. Anderson. 2010. “Procrastination and the extended will”. In: *The Thief of Time: Philosophical Essays on Procrastination*. Ed. by A. Chrisoula and W. M. D. New York, New York, USA: Oxford University Press.
- [226] Heath, T. and C. Bizer. 2011. “Linked Data: Evolving the Web into a Global Data Space”. *Synthesis Lectures on the Semantic Web: Theory and Technology*. 1(1): 1–136.
- [227] Hedström, P. and P. Ylikoski. 2010. “Causal mechanisms in the social sciences”. *Annual Review of Sociology*. 36: 49–67.
- [228] Heersmink, R. 2015. “Dimensions of integration in embedded and extended cognitive systems”. *Phenomenology and the Cognitive Sciences*. 14(3): 577–598.
- [229] Heersmink, R. 2016. “The Internet, Cognitive Enhancement, and the Values of Cognition”. *Minds and Machines*. 26(4): 389–407.

- [230] Heersmink, R. 2017a. “Distributed cognition and distributed morality: Agency, artifacts and systems”. *Science and Engineering Ethics*. 23(2): 431–448.
- [231] Heersmink, R. 2017b. “Extended mind and cognitive enhancement: Moral aspects of cognitive artifacts”. *Phenomenology and the Cognitive Sciences*. 16(1): 17–32.
- [232] Heersmink, R. in press. “Distributed selves: Personal identity and extended memory systems”. *Synthese*.
- [233] Heintz, C. 2006. “Web search engines and distributed assessment systems”. *Pragmatics & Cognition*. 14(2): 387–409.
- [234] Hendler, J. and A. M. Mulvehill. 2016. *Social Machines: The Coming Collision of Artificial Intelligence, Social Networking, and Humanity*. New York, New York, USA: Apress.
- [235] Hendler, J., N. R. Shadbolt, W. Hall, T. Berners-Lee, and D. Weitzner. 2008. “Web Science: An Interdisciplinary Approach to Understanding the Web”. *Communications of the ACM*. 51(7): 60–69.
- [236] Hendler, J. and T. Berners-Lee. 2010. “From the Semantic Web to social machines: A research challenge for AI on the World Wide Web”. *Artificial Intelligence*. 174: 156–161.
- [237] Herrmann, E., J. Call, M. V. Hernández-Lloreda, B. Hare, and M. Tomasello. 2007. “Humans have evolved specialized skills of social cognition: The cultural intelligence hypothesis”. *Science*. 317(5843): 1360–1366.
- [238] Heylighen, F. 2013. “From Human Computation to the Global Brain: The Self-Organization of Distributed Intelligence”. In: *Handbook of Human Computation*. Ed. by P. Michelucci. New York, New York, USA: Springer.
- [239] Heylighen, F. 2016. “Stigmergy as a universal coordination mechanism I: Definition and components”. *Cognitive Systems Research*. 38: 4–13.
- [240] Hollan, J., E. Hutchins, and D. Kirsh. 2000. “Distributed cognition: Toward a new foundation for human-computer interaction research”. *ACM Transactions on Computer-Human Interaction*. 7(2): 174–196.

- [241] Hollnagel, E. 2001. “Extended cognition and the future of ergonomics”. *Theoretical Issues in Ergonomics Science*. 2(3): 309–315.
- [242] Hongladarom, S. 2004. “Making information transparent as a means to close the global digital divide”. *Minds and Machines*. 14(1): 85–99.
- [243] Hooper, C., B. Bailey, H. Glaser, and J. Hendler. 2016. “Social machines in practice: Solutions, stakeholders and scopes”. In: *8th International ACM Web Science Conference*. Hannover, Germany.
- [244] Horsman, D. C. 2015. “Abstraction/Representation Theory for heterotic physical computing”. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*. 373: 20140224.
- [245] Horvitz, E. J. 2007. “Reflections on challenges and promises of mixed-initiative interaction”. *AI Magazine*. 28(2): 19–22.
- [246] Hsu, J. 2014. “IBM’s new brain”. *IEEE Spectrum*. 51(10): 17–19.
- [247] Huang, L., A. D. Galinsky, D. H. Gruenfeld, and L. E. Guillory. 2011. “Powerful postures versus powerful roles: Which is the proximate correlate of thought and behavior?” *Psychological Science*. 22(1): 95–102.
- [248] Huber, L. 2015. “Otto’s Notebook vs the Memex: Comparing the smartphone and the PC as two paradigms of cognitive enhancement”. *gnovis*. XV(2): 1–12.
- [249] Huberman, B. A., C. H. Loch, and A. Öngüler. 2004. “Status as a valued resource”. *Social Psychology Quarterly*. 67(1): 103–114.
- [250] Huberman, B. A., D. M. Romero, and F. Wu. 2009. “Crowdsourcing, attention and productivity”. *Journal of Information Science*. 35(6): 758–765.
- [251] Hull, G. 2016. “Apple, the FBI, iPhones, and the Extended Mind Hypothesis”. <http://www.newappsblog.com/2016/02/apple-the-fbi-iphones-and-the-extended-mind-hypothesis.html>. Accessed: 20th August 2016.
- [252] Hurley, S. 1998. *Consciousness in Action*. Cambridge, Massachusetts, USA: Harvard University Press.

- [253] Hutchins, E. 1991. “The social organization of distributed cognition”. In: *Perspectives on Socially Shared Cognition*. Ed. by L. Resnick, J. Levine, and S. Teasley. Washington DC, USA: The American Psychological Association.
- [254] Hutchins, E. 1995. *Cognition in the Wild*. Cambridge, Massachusetts, USA: MIT Press.
- [255] Hutchins, E. 2001. “Distributed Cognition”. In: *International Encyclopedia of the Social & Behavioral Sciences*. Ed. by N. J. Smelser and P. B. Baltes. Oxford, UK: Elsevier Science Ltd.
- [256] Hutchins, E. 2010. “Cognitive ecology”. *Topics in Cognitive Science*. 2(4): 705–715.
- [257] Hutchins, E. 2014. “The cultural ecosystem of human cognition”. *Philosophical Psychology*. 27(1): 34–49.
- [258] Hutchins, E., N. Weibel, C. Emmenegger, A. Fouse, and B. Holder. 2013. “An integrative approach to understanding flight crew activity”. *Journal of Cognitive Engineering and Decision Making*. 7(4): 353–379.
- [259] Iriki, A., M. Tanaka, and Y. Iwamura. 1996. “Coding of modified body schema during tool use by macaque postcentral neurones”. *Neuroreport*. 7(14): 2325–2330.
- [260] Ishimaru, S., T. Dingler, K. Kunze, K. Kise, and A. Dengel. 2016. “Reading interventions: Tracking reading state and designing interventions”. In: *ACM International Joint Conference on Pervasive and Ubiquitous Computing*. Heidelberg, Germany.
- [261] Jacucci, G., L. Gamberini, J. Freeman, and A. Spagnolli, eds. 2014. *Symbiotic Interaction*. London, UK: Springer.
- [262] Janssen, J. H., J. N. Bailenson, W. Ijsselstein, and J. H. D. M. Westerink. 2010. “Intimate heartbeats: Opportunities for affective communication technology”. *IEEE Transactions on Affective Computing*. 1(2): 72–80.
- [263] Johnson, C. R., E. Hendriks, I. J. Berezhnoy, E. Brevdo, S. M. Hughes, I. Daubechies, J. Li, E. Postma, and J. Z. Wang. 2008. “Image processing for artist identification”. *IEEE Signal Processing Magazine*. 25(4): 37–48.

- [264] Judd, S., M. Kearns, and Y. Vorobeychik. 2010. “Behavioral dynamics and influence in networked coloring and consensus”. *Proceedings of the National Academy of Sciences*. 107(34): 14978–14982.
- [265] Kaiser, M. I. in press. “The Components and Boundaries of Mechanisms”. In: *The Routledge Handbook of Mechanisms and Mechanical Philosophy*. Ed. by S. Glennan and P. M. Illari. New York, New York, USA: Routledge.
- [266] Kaiser, M. I. and B. Krickel. in press. “The metaphysics of constitutive mechanistic phenomena”. *The British Journal for the Philosophy of Science*.
- [267] Kamitani, Y. and F. Tong. 2005. “Decoding the visual and subjective contents of the human brain”. *Nature Neuroscience*. 8(5): 679–685.
- [268] Kaplan, D. M. 2012. “How to demarcate the boundaries of cognition”. *Biology & Philosophy*. 27(4): 545–570.
- [269] Kapoor, A., D. Tan, P. Shenoy, and E. Horvitz. 2008. “Complementary computing for visual tasks: Meshing computer vision with human visual processing”. In: *8th IEEE International Conference on Automatic Face & Gesture Recognition*. Amsterdam, The Netherlands.
- [270] Kearns, M. 2012. “Experiments in social computation”. *Communications of the ACM*. 55(10): 56–67.
- [271] Kendon, V., A. Sebald, and S. Stepney. 2015. “Heterotic computing: Past, present and future”. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*. 373: 20140225.
- [272] Kennedy, J., R. C. Eberhart, and Y. Shi. 2001. *Swarm Intelligence*. San Francisco, California, USA: Morgan Kaufmann Publishers.
- [273] Khatib, F., S. Cooper, M. D. Tyka, K. Xu, I. Makedon, Z. Popović, D. Baker, and Foldit Players. 2011a. “Algorithm discovery by protein folding game players”. *Proceedings of the National Academy of Sciences*. 108(47): 18949–18953.

- [274] Khatib, F., F. DiMaio, S. Cooper, M. Kazmierczyk, M. Gilski, S. Krzywda, H. Zabranska, I. Pichova, J. Thompson, Z. Popović, M. Jaskolski, and D. Baker. 2011b. “Crystal structure of a monomeric retroviral protease solved by protein folding game players”. *Nature Structural & Molecular Biology*. 18(10): 1175–1177.
- [275] Kiourti, A., C. Lee, and J. L. Volakis. 2016. “Fabrication of textile antennas and circuits with 0.1 mm precision”. *IEEE Antennas and Wireless Propagation Letters*. 15: 151–153.
- [276] Kirsh, D. 2013. “Embodied cognition and the magical future of interaction design”. *ACM Transactions on Computer-Human Interaction*. 20(1): 1–30.
- [277] Kirsh, D. and P. Maglio. 1994. “On distinguishing epistemic from pragmatic action”. *Cognitive Science*. 18: 513–549.
- [278] Klein, G., S. Wiggins, and C. O. Dominguez. 2010. “Team sense-making”. *Theoretical Issues in Ergonomics Science*. 11(4): 304–320.
- [279] Klemmer, S. R., B. Hartmann, and L. Takayama. 2006. “How bodies matter: Five themes for interaction design”. In: *6th Conference on Designing Interactive Systems*. University Park, Pennsylvania, USA.
- [280] Kohl, N. E., E. A. Emini, W. A. Schleif, L. J. Davis, J. C. Heimbach, R. A. F. Dixon, E. M. Scolnick, and I. S. Sigal. 1988. “Active human immunodeficiency virus protease is required for viral infectivity”. *Proceedings of the National Academy of Sciences*. 85(13): 4686–4690.
- [281] Konrad, A., E. Isaacs, and S. Whittaker. 2016. “Technology Mediated Memory: Is Technology Altering Our Memories And Interfering With Well-Being?” *ACM Transactions on Computer-Human Interaction*. 24(4): 1–29.
- [282] Koriat, A. 2000. “The feeling of knowing: Some metatheoretical implications for consciousness and control”. *Consciousness and Cognition*. 9(2): 149–171.

- [283] Koriat, A. and R. Nussinson. 2009. “Attributing study effort to data-driven and goal-driven effects: Implications for metacognitive judgments”. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 35(5): 1338–1343.
- [284] Krafft, P., K. Zhou, I. Edwards, K. Starbird, and E. Spiro. 2017. “Centralized, parallel, and distributed information processing during collective sensemaking”. In: *ACM Conference on Human Factors in Computing Systems (CHI’17)*. Denver, Colorado, USA.
- [285] Krötzsch, M., D. Vrandečić, M. Völkel, H. Haller, and R. Studer. 2007. “Semantic Wikipedia”. *Journal of Web Semantics*. 5(4): 251–261.
- [286] Krueger, J. 2012. “Seeing mind in action”. *Phenomenology and the Cognitive Sciences*. 11(2): 149–173.
- [287] Kummer, H., L. Daston, G. Gigerenzer, and J. B. Silk. 1997. “The social intelligence hypothesis”. In: *Human By Nature: Between Biology and the Social Sciences*. Ed. by P. Weingart, S. D. Mitchell, P. J. Richerson, and S. Maasen. Mahwah, New Jersey, USA: Lawrence Erlbaum Associates.
- [288] Kunda, Z. 1999. *Social Cognition: Making Sense of People*. Cambridge, Massachusetts, USA: MIT Press.
- [289] Kunze, K., M. Iwamura, K. Kise, S. Uchida, and S. Omachi. 2013. “Activity Recognition for the Mind: Toward a Cognitive “Quantified Self””. *Computer*. 46(10): 105–108.
- [290] Kurs, A., A. Karalis, R. Moffatt, J. D. Joannopoulos, P. Fisher, and M. Soljačić. 2007. “Wireless Power Transfer via Strongly Coupled Magnetic Resonances”. *Science*. 317(5834): 83–86.
- [291] Kurzhals, K. and D. Weiskopf. 2015. “Eye Tracking for Personal Visual Analytics”. *IEEE Computer Graphics and Applications*. 35(4): 64–72.
- [292] Lakoff, G. and M. Johnson. 1980. *Metaphors We Live By*. Chicago, Illinois, USA: University of Chicago Press.
- [293] Laland, K. N. and M. J. O’Brien. 2011. “Cultural niche construction: An introduction”. *Biological Theory*. 6(3): 191–202.

- [294] Laland, K. N., J. Odling-Smee, and M. W. Feldman. 2000. "Niche construction, biological evolution, and cultural change". *Behavioral and Brain Sciences*. 23(1): 131–146.
- [295] Lampel, J. and A. Bhalla. 2007. "The role of status seeking in online communities: Giving the gift of experience". *Journal of Computer-Mediated Communication*. 12(2): 434–455.
- [296] Lampos, V., T. De Bie, and N. Cristianini. 2010. "Flu detector—tracking epidemics on Twitter". In: *European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases*. Barcelona, Spain.
- [297] Landes, D. S. 2000. *Revolution in Time: Clocks and the Making of the Modern World*. London, UK: Viking Press.
- [298] Lanzing, M. 2016. "The transparent self". *Ethics and Information Technology*. 18(1): 9–16.
- [299] Lasecki, W. S. and J. P. Bigham. 2013. "Interactive Crowds: Real-Time Crowdsourcing and Crowd Agents". In: *Handbook of Human Computation*. Ed. by P. Michelucci. New York, New York, USA: Springer.
- [300] Law, E. and L. von Ahn. 2011. "Human Computation". *Synthesis Lectures on Artificial Intelligence and Machine Learning*. 5(3): 1–121.
- [301] Lazer, D. and A. Friedman. 2007. "The Network Structure of Exploration and Exploitation". *Administrative Science Quarterly*. 52(4): 667–694.
- [302] Lazer, D., A. S. Pentland, L. Adamic, S. Aral, A.-L. Barabási, D. Brewer, N. Christakis, N. Contractor, J. Fowler, and M. Gutmann. 2009. "Life in the network: The coming age of computational social science". *Science*. 323(5915): 721–723.
- [303] Leavitt, H. J. 1951. "Some effects of certain communication patterns on group performance". *The Journal of Abnormal and Social Psychology*. 46(1): 38–50.
- [304] Levy, N. 2007. *Neuroethics: Challenges for the 21st Century*. Cambridge, UK: Cambridge University Press.

- [305] Lindblom, J. and T. Ziemke. 2003. “Social Situatedness of Natural and Artificial Intelligence: Vygotsky and Beyond”. *Adaptive Behavior*. 11(2): 79–96.
- [306] Linden, G., B. Smith, and J. York. 2003. “Amazon.com Recommendations: Item-to-Item Collaborative Filtering”. *IEEE Internet Computing*. 7(1): 76–80.
- [307] Lingley, A. R., M. Ali, Y. Liao, R. Mirjalili, M. Klonner, M. Sopenan, S. Suihkonen, T. Shen, B. P. Otis, H. Lipsanen, and B. A. Parviz. 2011. “A single-pixel wireless contact lens display”. *Journal of Micromechanics and Microengineering*. 21(12): 125014.
- [308] Lintott, C. J. and J. Reed. 2013. “Human Computation in Citizen Science”. In: *Handbook of Human Computation*. Ed. by P. Michelucci. New York, New York, USA: Springer.
- [309] Lintott, C. J., K. Schawinski, W. Keel, H. Van Arkel, N. Bennert, E. Edmondson, D. Thomas, D. J. B. Smith, P. D. Herbert, M. J. Jarvis, S. Virani, D. Andreescu, S. P. Bamford, K. Land, P. Murray, R. C. Nichol, M. J. Raddick, A. Slosar, A. Szalay, and J. Vandenberg. 2009. “Galaxy Zoo: ‘Hanny’s Voorwerp’, a quasar light echo?” *Monthly Notices of the Royal Astronomical Society*. 399(1): 129–140.
- [310] Lintott, C. J., K. Schawinski, A. Slosar, K. Land, S. Bamford, D. Thomas, M. J. Raddick, R. C. Nichol, A. Szalay, D. Andreescu, P. Murray, and J. van den Berg. 2008. “Galaxy Zoo: Morphologies derived from visual inspection of galaxies from the Sloan Digital Sky Survey”. *Monthly Notices of the Royal Astronomical Society*. 389(3): 1179–1189.
- [311] Liu, V., A. Parks, V. Talla, S. Gollakota, D. Wetherall, and J. R. Smith. 2013. “Ambient backscatter: Wireless communication out of thin air”. *ACM SIGCOMM Computer Communication Review*. 43(4): 39–50.
- [312] Loh, K. K. and R. Kanai. 2016. “How has the Internet reshaped human cognition?” *The Neuroscientist*. 22(5): 506–520.

- [313] Lorenz, J., H. Rauhut, F. Schweitzer, and D. Helbing. 2011. “How social influence can undermine the wisdom of crowd effect”. *Proceedings of the National Academy of Sciences*. 108(22): 9020–9025.
- [314] Ludwig, D. 2015. “Extended cognition and the explosion of knowledge”. *Philosophical Psychology*. 28(3): 355–368.
- [315] Lupton, D. 2013. “Understanding the Human Machine”. *IEEE Technology and Society Magazine*. 32(4): 25–30.
- [316] Lupton, D. 2014. “Self-tracking cultures: Towards a sociology of personal informatics”. In: *26th Australian Computer-Human Interaction Conference*. Sydney, Australia.
- [317] Lupton, D. 2015. *Digital Sociology*. Abingdon, England, UK: Routledge.
- [318] Lynch, M. P. 2014. “Neuromedia, extended knowledge and understanding”. *Philosophical Issues*. 24(1): 299–313.
- [319] Mace, W. M. 1977. “James J. Gibson’s strategy for perceiving: Ask not what’s inside your head, but what your head’s inside of”. In: *Perceiving, Acting, and Knowing*. Ed. by R. Shaw and J. Bransford. Hillsdale, New Jersey, USA: Lawrence Erlbaum Associates.
- [320] Magid, M. and J. Reichenberg. 2015. “Botulinum toxin for depression? An idea that’s raising some eyebrows”. *Current Psychiatry*. 14(11): 43–56.
- [321] Malafouris, L. 2013. *How Things Shape the Mind: A Theory of Material Engagement*. Cambridge, Massachusetts, USA: MIT Press.
- [322] Malone, T. W. and M. S. Bernstein, eds. 2015. *Handbook of Collective Intelligence*. Cambridge, Massachusetts, USA: MIT Press.
- [323] Malone, T. W., R. Laubacher, and C. Dellarocas. 2010. “The Collective Intelligence Genome”. *MIT Sloan Management Review*. 51(3): 21–31.
- [324] Manasian, D. 2003. “Digital dilemmas: A survey of the Internet society”. *The Economist*. 366(8308).

- [325] Mann, S. 1997. “Wearable computing: A first step toward personal imaging”. *Computer*. 30(2): 25–32.
- [326] Mansour, O. 2009. “Group Intelligence: A distributed cognition perspective”. In: *International Conference on Intelligent Networking and Collaborative Systems*. Barcelona, Spain.
- [327] Mao, A., W. Mason, S. Suri, and D. J. Watts. 2016. “An experimental study of team size and performance on a complex task”. *PLoS One*. 11(4): e0153048.
- [328] Mason, W. 2013. “Collective Search as Human Computation”. In: *Handbook of Human Computation*. Ed. by P. Michelucci. New York, New York, USA: Springer.
- [329] Mason, W. A., F. R. Conrey, and E. R. Smith. 2007. “Situating social influence processes: Dynamic, multidirectional flows of influence within social networks”. *Personality and Social Psychology Review*. 11(3): 279–300.
- [330] Mason, W. A., A. Jones, and R. L. Goldstone. 2008. “Propagation of innovations in networked groups”. *Journal of Experimental Psychology: General*. 137(3): 422–433.
- [331] Mason, W. and D. J. Watts. 2012. “Collaborative learning in networks”. *Proceedings of the National Academy of Sciences*. 109(3): 764–769.
- [332] Matsumoto, T., S. Hashimoto, and N. Okude. 2008. “The Embodied Web: Embodied Web-services interaction with an umbrella for augmented city experiences”. *Computer Animation and Virtual Worlds*. 19(1): 49–66.
- [333] Maurer, H. 2015. “Does the Internet make us stupid?” *Communications of the ACM*. 58(1): 48–51. 2629544.
- [334] Mayer-Schönberger, V. 2007. “Useful void: The art of forgetting in the age of ubiquitous computing”. *Tech. rep.* No. KSG Working Paper RWP07-022. Oxford Internet Institute, University of Oxford, Oxford, UK.
- [335] Mayer-Schönberger, V. 2011. *Delete: The Virtue of Forgetting in the Digital Age*. Princeton, New Jersey, USA: Princeton University Press.

- [336] McBeath, M. K., D. M. Shaffer, and M. K. Kaiser. 1995. “How baseball outfielders determine where to run to catch fly balls”. *Science*. 268(5210): 569–573.
- [337] McDonald, A., C. Haslam, P. Yates, B. Gurr, G. Leeder, and A. Sayers. 2011. “Google Calendar: A new memory aid to compensate for prospective memory deficits following acquired brain injury”. *Neuropsychological Rehabilitation*. 21(6): 784–807.
- [338] Meira, S. R. L., V. A. A. Burégio, L. M. Nascimento, E. Figueiredo, M. Neto, B. Encarnação, and V. C. Garcia. 2011. “The Emerging Web of Social Machines”. In: *35th Annual Computer Software and Applications Conference (COMPSAC)*. Munich, Germany.
- [339] Mejova, Y., I. Weber, and M. W. Macy, eds. 2015. *Twitter: A Digital Socioscope*. New York, New York, USA: Cambridge University Press.
- [340] Menary, R., ed. 2010. *The Extended Mind*. Cambridge, Massachusetts, USA: MIT Press.
- [341] Mercier, H. and D. Sperber. 2011. “Why do humans reason? Arguments for an argumentative theory”. *Behavioral and Brain Sciences*. 34(2): 57–111.
- [342] Merolla, P. A., J. V. Arthur, R. Alvarez-Icaza, A. S. Cassidy, J. Sawada, F. Akopyan, B. L. Jackson, N. Imam, C. Guo, and Y. Nakamura. 2014. “A million spiking-neuron integrated circuit with a scalable communication network and interface”. *Science*. 345(6197): 668–673.
- [343] Meshi, D., D. I. Tamir, and H. R. Heekeren. 2015. “The emerging neuroscience of social media”. *Trends in Cognitive Sciences*. 19(12): 771–782.
- [344] Metzger, M. J. 2007. “Making sense of credibility on the Web: Models for evaluating online information and recommendations for future research”. *Journal of the American Society for Information Science and Technology*. 58(13): 2078–2091.
- [345] Meyer, E. T. and R. Schroeder. 2015. *Knowledge Machines: Digital Transformations of the Sciences and Humanities*. Cambridge, Massachusetts, USA: MIT Press.

- [346] Michaelian, K. 2012. “Is external memory memory? Biological memory and extended mind”. *Consciousness and Cognition*. 21(3): 1154–1165.
- [347] Michaelian, K. and S. Arango-Muñoz. 2014. “Epistemic feelings, epistemic emotions: Review and introduction to the focus section”. *Philosophical Inquiries*. 2(1): 97–122.
- [348] Michelucci, P., ed. 2013. *Handbook of Human Computation*. New York, New York, USA: Springer.
- [349] Michelucci, P. 2016. “Human computation and convergence”. In: *Handbook of Science and Technology Convergence*. Ed. by W. S. Bainbridge and M. C. Roco. New York, New York, USA: Springer.
- [350] Michelucci, P. and J. L. Dickinson. 2016. “The power of crowds”. *Science*. 351(6268): 32–33.
- [351] Miller, B. and I. Record. 2013. “Justified belief in a digital age: On the epistemic implications of secret Internet technologies”. *Episteme*. 10(02): 117–134.
- [352] Miller, G. A. 2003. “The cognitive revolution: A historical perspective”. *Trends in Cognitive Sciences*. 7(3): 141–144.
- [353] Miller, J. 2014. “The fourth screen: Mediatization and the smart-phone”. *Mobile Media & Communication*. 2(2): 209–226.
- [354] Miller, M. 2015. *The Internet of Things: How Smart TVs, Smart Cars, Smart Homes, and Smart Cities are Changing the World*. Indianapolis, Indiana, USA: Que Publishing.
- [355] Mills, K. L. 2014. “Effects of Internet use on the adolescent brain: Despite popular claims, experimental evidence remains scarce”. *Trends in Cognitive Sciences*. 18(8): 385–387.
- [356] Mills, K. L. 2016. “Possible effects of Internet use on cognitive development in adolescence”. *Media and Communication*. 4(3): 4–12.
- [357] Montaña, F. G. and J. Vallverdú. 2015. “Hackable Bodies. An Unknown Error Has Occurred”. *Journal of the New Media Caucus*. 11(2).

- [358] Moore, R. 2011. "Digital file sharing: An examination of neutralization and rationalization techniques employed by digital file sharers". In: *Cyber Criminology: Exploring Internet Crimes and Criminal Behavior*. Ed. by K. Jaishankar. Boca Raton, Florida, USA: CRC Press.
- [359] Morgan, J. 2016. "Gaming for dementia research: A quest to save the brain". *The Lancet Neurology*. 16(13): 1313.
- [360] Negri, P., L. Gamberini, and S. Cutini. 2014. "A Review of the Research on Subliminal Techniques for Implicit Interaction in Symbiotic Systems". In: *Symbiotic Interaction*. Ed. by G. Jacucci, L. Gamberini, J. Freeman, and A. Spagnolli. *Lecture Notes in Computer Science*. London, UK: Springer.
- [361] Neisser, U. 1997. "The ecological study of memory". *Philosophical Transactions of the Royal Society B: Biological Sciences*. 352(1362): 1697–1701.
- [362] Newman, E. A. 2003. "New roles for astrocytes: Regulation of synaptic transmission". *Trends in Neurosciences*. 26(10): 536–542.
- [363] Nhan, J., L. Huey, and R. Broll. 2017. "Digilantism: An analysis of crowdsourcing and the Boston marathon bombings". *British Journal of Criminology*. 57(2): 341–361.
- [364] Norman, D. A. 1992. "Design principles for cognitive artifacts". *Research in Engineering Design*. 4(1): 43–50.
- [365] Norman, D. A. 1993. *Things That Makes Us Smart*. Reading, Massachusetts, USA: Addison-Wesley.
- [366] Odling-Smee, J., K. N. Laland, and M. W. Feldman. 2003. *Niche Construction: The Neglected Process in Evolution*. Princeton, New Jersey, USA: Princeton University Press.
- [367] O'Hara, K. 2013. "The Technology of Collective Memory and the Normativity of Truth". In: *Philosophy and Engineering: Reflections on Practice, Principles and Process*. Ed. by D. P. Michelfelder, N. McCarthy, and D. E. Goldberg. London, UK: Springer.

- [368] O'Hara, K., N. Contractor, W. Hall, J. A. Hendler, and N. Shadbolt. 2013. "Web Science: Understanding the Emergence of Macro-Level Features on the World Wide Web". *Foundations and Trends in Web Science*. 4(2–3): 103–267.
- [369] O'Hara, K. and N. R. Shadbolt. 2008. *The Spy in the Coffee Machine: The End of Privacy as We Know It*. Oxford, UK: Oneworld Publications.
- [370] O'Hara, K., M. Tuffield, and N. Shadbolt. 2008. "Lifelogging: Privacy and empowerment with Memories for Life". *Identity in the Information Society*. 1(1): 155–172.
- [371] Oksman, V. and P. Rautiainen. 2003a. "Extension of the hand: Children's and teenagers' relationship with the mobile phone in Finland". In: *Mediating the Human Body: Technology, Communication and Fashion*. Ed. by L. Fortunati, J. E. Katz, and R. Riccini. Mahwah, New Jersey, USA: Lawrence Erlbaum.
- [372] Oksman, V. and P. Rautiainen. 2003b. "Perhaps it is a body part": How the mobile phone became an organic part of the everyday lives of Finnish children and teenagers". In: *Machines that Become Us: The Social Context of Communication Technology*. Ed. by J. E. Katz. New Brunswick, New Jersey, USA: Transaction Publishers.
- [373] Oliveira, J. F., V. M. Sardinha, S. Guerra-Gomes, A. Araque, and N. Sousa. 2015. "Do stars govern our actions? Astrocyte involvement in rodent behavior". *Trends in Neurosciences*. 38(9): 535–549.
- [374] Ong, W. J. 1982. *Orality and Literacy: The Technologizing of the Word*. London, UK: Methuen & Co. Ltd.
- [375] Oppenheimer, D. M. 2008. "The secret life of fluency". *Trends in Cognitive Sciences*. 12(6): 237–241.
- [376] Orwell, G. 2004. *Nineteen Eighty-Four*. London, UK: Penguin Books Ltd.
- [377] Ottati, V., E. D. Price, C. Wilson, and N. Sumaktoyo. 2015. "When self-perceptions of expertise increase closed-minded cognition: The earned dogmatism effect". *Journal of Experimental Social Psychology*. 61: 131–138.

- [378] Paglieri, F. 2012. "Ulysses' will: Self control, external constraints, and games". In: *Consciousness in Interaction: The Role of the Natural and Social Context in Shaping Consciousness*. Ed. by F. Paglieri. Amsterdam, The Netherlands: John Benjamins Publishing Company.
- [379] Palermos, S. O. 2011. "Belief-forming processes, extended". *Review of Philosophy and Psychology*. 2(4): 741–765.
- [380] Palermos, S. O. 2015. "Active externalism, virtue reliabilism and scientific knowledge". *Synthese*. 192(9): 2955–2986.
- [381] Palermos, S. O. 2016. "The dynamics of group cognition". *Minds and Machines*. 26(4): 409–440.
- [382] Palermos, S. O. in press. "Social Machines: A Philosophical Engineering". *Phenomenology and the Cognitive Sciences*.
- [383] Paradiso, J. and T. Starner. 2005. "Energy scavenging for mobile and wireless electronics". *IEEE Pervasive Computing*. 4(1): 18–27.
- [384] Parkinson, B. 2015. "Social perception and attribution". In: *An Introduction to Social Psychology*. Ed. by M. Hewstone, W. Stroebe, and K. Jonas. 6th. Chichester, West Sussex, UK: John Wiley & Sons Ltd.
- [385] Parunak, H. V. D. 2005. "A survey of environments and mechanisms for human-human stigmergy". In: *International Workshop on Environments for Multi-Agent Systems*. Utrecht, The Netherlands.
- [386] Pecher, D. and R. Zwaan, eds. 2005. *Grounding Cognition: The Role of Perception and Action in Memory, Language, and Thinking*. Cambridge, UK: Cambridge University Press.
- [387] Perea, G., M. Sur, and A. Araque. 2014. "Neuron-glia networks: Integral gear of brain function". *Frontiers in Cellular Neuroscience*. 8(378): 1–9.
- [388] Ping, R. and S. Goldin-Meadow. 2010. "Gesturing saves cognitive resources when talking about nonpresent objects". *Cognitive Science*. 34(4): 602–619.

- [389] Planqué, R., J. B. Van Den Berg, and N. R. Franks. 2010. "Recruitment strategies and colony size in ants". *PLoS One*. 5(8): e11664.
- [390] Pogue, D. 2017. "How to stamp out fake news". *Scientific American*. 316(2): 20.
- [391] Prensky, M. 2001. "Digital Natives, Digital Immigrants, Part II: Do They Really Think Differently". *On The Horizon*. 9(6): 1–6.
- [392] Pritchard, D. 2009. *Knowledge*. Basingstoke, England, UK: Palgrave Macmillan.
- [393] Pritchard, D. 2010. "Cognitive ability and the extended cognition thesis". *Synthese*. 175(1): 133–151.
- [394] Pritchard, D. 2012. "Anti-luck virtue epistemology". *The Journal of Philosophy*. 109(3): 247–279.
- [395] Pritchard, D. 2014. "Virtue Epistemology, Extended Cognition, and the Epistemology of Education". *Universitas: Monthly Review of Philosophy and Culture*. 478: 47–66.
- [396] Proust, J. 2013. *The Philosophy of Metacognition: Mental Agency and Self-Awareness*. Oxford, UK: Oxford University Press.
- [397] Quercia, D., M. Kosinski, D. Stillwell, and J. Crowcroft. 2011. "Our Twitter profiles, our selves: Predicting personality with Twitter". In: *IEEE Third International Conference on Social Computing (SocialCom)*. Boston, Massachusetts, USA.
- [398] Quinn, A. J. and B. B. Bederson. 2011. "Human computation: A survey and taxonomy of a growing field". In: *Conference on Human Factors in Computing Systems*. Vancouver, British Columbia, Canada.
- [399] Rademaker, R. L., D.-A. Wu, I. M. Bloem, and A. T. Sack. 2014. "Intensive tool-practice and skillfulness facilitate the extension of body representations in humans". *Neuropsychologia*. 56: 196–203.
- [400] Rainie, L. and B. Wellman. 2012. *Networked: The New Social Operating System*. Cambridge, Massachusetts, USA: MIT Press.

- [401] Randić, S., R. C. Connor, W. B. Sherwin, and W. Krützen. 2012. “A novel mammalian social structure in Indo-Pacific bottlenose dolphins (*Tursiops* sp.): Complex male alliances in an open social network”. *Proceedings of the Royal Society B*. 279(1740): 3083–3090.
- [402] Ratkiewicz, J., M. Conover, M. R. Meiss, B. Gonçalves, A. Flammini, and F. Menczer. 2011. “Detecting and Tracking Political Abuse in Social Media”. In: *5th International AAAI Conference on Weblogs and Social Media*. Barcelona, Spain.
- [403] Record, I. and B. Miller. in press. “Taking iPhone Seriously: Epistemic Technologies and the Extended Mind”. In: *Extended Epistemology*. Ed. by A. J. Carter, A. Clark, J. Kallestrup, O. S. Palermos, and D. Pritchard. Oxford, UK: Oxford University Press.
- [404] Reiner, P. B. and S. K. Nagel. in press. “Technologies of the extended mind: Defining the issues”. In: *Neuroethics: Anticipating the Future*. Ed. by J. Illes and S. Hossain. Oxford, UK: Oxford University Press.
- [405] Rheingold, H. 2012. *Net Smart: How to Thrive Online*. Cambridge, Massachusetts, USA: MIT Press.
- [406] Ricci, A., M. Piunti, L. Tummolini, and C. Castelfranchi. 2015. “The Mirror World: Preparing for Mixed-Reality Living”. *IEEE Pervasive Computing*. 14(2): 60–63.
- [407] Richardson, I. 2007. “Pocket technospaces: The bodily incorporation of mobile media”. *Continuum: Journal of Media & Cultural Studies*. 21(2): 205–215.
- [408] Risko, E. F. and S. J. Gilbert. 2016. “Cognitive Offloading”. *Trends in Cognitive Sciences*. 20(9): 676–688.
- [409] Robbins, P. and M. Aydede, eds. 2009. *The Cambridge Handbook of Situated Cognition*. New York, New York, USA: Cambridge University Press.
- [410] Robertson, D. and F. Giunchiglia. 2013. “Programming the social computer”. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*. 371(1987): 20120379.

- [411] Rosenquist, J. N., J. H. Fowler, and N. A. Christakis. 2011. “Social network determinants of depression”. *Molecular Psychiatry*. 16(3): 273–281.
- [412] Rowlands, M. 1999. *The Body in Mind*. New York, New York, USA: Cambridge University Press.
- [413] Rowlands, M. 2006. *Body Language: Representation in Action*. Cambridge, Massachusetts, USA: MIT Press.
- [414] Rowlands, M. 2010. *The New Science of the Mind: From Extended Mind to Embodied Phenomenology*. Cambridge, Massachusetts, USA: MIT Press.
- [415] Ruckenstein, M. 2014. “Visualized and interacted life: Personal analytics and engagements with data doubles”. *Societies*. 4(1): 68–84.
- [416] Ruiter, J. de, G. Weston, and S. M. Lyon. 2011. “Dunbar’s number: Group size and brain physiology in humans reexamined”. *American Anthropologist*. 113(4): 557–568.
- [417] Rupert, R. D. 2004. “Challenges to the Hypothesis of Extended Cognition”. *Journal of Philosophy*. 101(8): 389–428.
- [418] Rupert, R. D. 2009. *Cognitive Systems and the Extended Mind*. New York, New York, USA: Oxford University Press.
- [419] Sakaki, T., M. Okazaki, and Y. Matsuo. 2010. “Earthquake shakes Twitter users: Real-time event detection by social sensors”. In: *19th International World Wide Web Conference*. Raleigh, North Carolina, USA.
- [420] Saleh, B. and A. Elgammal. 2016. “Large-scale classification of fine-art paintings: Learning the right metric on the right feature”. *International Journal for Digital Art History*. (2): 71–93.
- [421] Salerno, S., E. Zamagni, C. Urquizar, R. Salemme, A. Farnè, and F. Frassinetti. 2012. “Increases of corticospinal excitability in self-related processing”. *European Journal of Neuroscience*. 36(5): 2716–2721.
- [422] Salmon, W. C. 1984. *Scientific Explanation and the Causal Structure of the World*. Princeton, New Jersey, USA: Princeton University Press.

- [423] Schechtman, M. 2011. “The narrative self”. In: *The Oxford Handbook of the Self*. Ed. by S. Gallagher. Oxford, UK: Oxford University Press.
- [424] schraefel, m., D. A. Smith, A. Owens, A. Russell, C. Harris, and M. Wilson. 2005. “The evolving mSpace platform: Leveraging the Semantic Web on the trail of the memex”. In: *Hypertext 2005 Conference*. Salzburg, Austria.
- [425] Schwamb, M. E., J. A. Orosz, J. A. Carter, W. F. Welsh, D. A. Fischer, G. Torres, A. W. Howard, J. R. Crepp, W. C. Keel, and C. J. Lintott. 2013. “Planet Hunters: A transiting circumbinary planet in a quadruple star system”. *The Astrophysical Journal*. 768(127): 1–21.
- [426] Segal, A. 2011. “Do Gestural Interfaces Promote Thinking? Embodied Interaction: Congruent Gestures and Direct-Touch Promote Performance in Math”. *Ph.D.* New York, New York, USA: Columbia University.
- [427] Shadbolt, N. R. 2013. “Knowledge acquisition and the rise of social machines”. *International Journal of Human-Computer Studies*. 71(2): 200–205.
- [428] Shadbolt, N. R. and T. Berners-Lee. 2008. “Web Science Emerges”. *Scientific American*. 299(4): 76–81.
- [429] Shadbolt, N. R. and P. R. Smart. 2015. “Knowledge Elicitation: Methods, Tools and Techniques”. In: *Evaluation of Human Work*. Ed. by J. R. Wilson and S. Sharples. 4th. Boca Raton, Florida, USA: CRC Press.
- [430] Shadbolt, N., W. Hall, and T. Berners-Lee. 2006. “The Semantic Web revisited”. *IEEE Intelligent Systems*. 21(3): 96–101.
- [431] Shadbolt, N., M. Van Kleek, and R. Binns. 2016. “The Rise of Social Machines: The development of a human/digital ecosystem”. *IEEE Consumer Electronics Magazine*. 5(2): 106–111.
- [432] Shapiro, L. 2013. “When is cognition embodied?” In: *Current Controversies in Philosophy of Mind*. Ed. by U. Kriegel. New York, New York, USA: Routledge.
- [433] Shapiro, L. A. 2007. “The embodied cognition research programme”. *Philosophy Compass*. 2(2): 338–346.

- [434] Shapiro, L. A. 2011. *Embodied Cognition*. Abingdon, Oxon, UK: Routledge.
- [435] Shapiro, L. A., ed. 2014. *The Routledge Handbook of Embodied Cognition*. New York, New York, USA: Routledge.
- [436] Shaw, M. E. 1964. “Communication Networks”. In: *Advances in Experimental Social Psychology*. Ed. by L. Berkowitz. Vol. 1. New York, New York, USA: Academic Press.
- [437] Silva, A. R., M. S. Pinho, L. Macedo, C. Moulin, S. Caldeira, and H. Firmino. 2017. “It is not only memory: Effects of sensecam on improving well-being in patients with mild alzheimer disease”. *International Psychogeriatrics*. 29(5): 741–754.
- [438] Simperl, E. and M. Luczak-Rösch. 2014. “Collaborative ontology engineering: A survey”. *The Knowledge Engineering Review*. 29(01): 101–131.
- [439] Simpson, T. W. 2012. “Evaluating Google as an epistemic tool”. *Metaphilosophy*. 43(4): 426–445.
- [440] Sims, K. 1994. “Evolving 3D morphology and behavior by competition”. *Artificial Life*. 1(4): 353–372.
- [441] Siorpaes, K. and M. Hepp. 2008. “Games with a purpose for the Semantic Web”. *IEEE Intelligent Systems*. 23(3): 50–60.
- [442] Siorpaes, K. and E. Simperl. 2010. “Human intelligence in the process of semantic content creation”. *World Wide Web*. 13(1–2): 33–59.
- [443] Small, G. W., T. D. Moody, P. Siddarth, and S. Y. Bookheimer. 2009. “Your brain on Google: Patterns of cerebral activation during Internet searching”. *American Journal of Geriatric Psychiatry*. 17(2): 116–126.
- [444] Smart, P. R. 2010a. “Cognition and the Web”. In: *1st ITA Workshop on Network-Enabled Cognition: The Contribution of Social and Technological Networks to Human Cognition*. Maryland, USA.

- [445] Smart, P. R. 2010b. “Extended Memory, the Extended Mind, and the Nature of Technology-Mediated Memory Enhancement”. In: *1st ITA Workshop on Network-Enabled Cognition: The Contribution of Social and Technological Networks to Human Cognition*. Maryland, USA.
- [446] Smart, P. R. 2012a. “Studying the Impact of Minority Views in a Computational Model of Collective Sensemaking: The Role of Network Structure”. In: *4th International Conference on Advanced Cognitive Technologies and Applications (COGNITIVE’12)*. Nice, France.
- [447] Smart, P. R. 2012b. “The Web-Extended Mind”. *Metaphilosophy*. 43(4): 426–445.
- [448] Smart, P. R. 2013. “Understanding the Cognitive Impact of Emerging Web Technologies: A Research Focus Area for Embodied, Extended and Distributed Approaches to Cognition”. In: *1st International Web for Wellbeing & Human Performance Workshop*. Paris, France.
- [449] Smart, P. R. 2014. “Embodiment, Cognition and the World Wide Web”. In: *The Routledge Handbook of Embodied Cognition*. Ed. by L. A. Shapiro. New York, New York, USA: Routledge.
- [450] Smart, P. R. 2017a. “COGWEB—Cognition and the Web: Social Minds, Distributed Brains, and Talkative Technologies”. In: *9th International Conference on Advanced Cognitive Technologies and Applications (COGNITIVE’17)*. Athens, Greece.
- [451] Smart, P. R. 2017b. “Machine Intelligence and the Social Web: How to Get a Cognitive Upgrade”. In: *9th International Conference on Advanced Cognitive Technologies and Applications (COGNITIVE’17)*. Athens, Greece.
- [452] Smart, P. R. 2017c. “Situating Machine Intelligence within the Cognitive Ecology of the Internet”. *Minds and Machines*. 27(2): 357–380.

- [453] Smart, P. R. in press(a). “Emerging Digital Technologies: Implications for Extended Conceptions of Cognition and Knowledge”. In: *Extended Epistemology*. Ed. by A. J. Carter, A. Clark, J. Kallestrup, O. S. Palermos, and D. Pritchard. Oxford, UK: Oxford University Press.
- [454] Smart, P. R. in press(b). “Extended Cognition and the Internet: A Review of Current Issues and Controversies”. *Philosophy & Technology*.
- [455] Smart, P. R. in press(c). “Mandevillian intelligence”. *Synthese*.
- [456] Smart, P. R. in press(d). “Mandevillian Intelligence: From Individual Vice to Collective Virtue”. In: *Socially-Extended Epistemology*. Ed. by A. J. Carter, A. Clark, J. Kallestrup, O. S. Palermos, and D. Pritchard. Oxford, UK: Oxford University Press.
- [457] Smart, P. R. Submitted(a). “Human-Extended Machine Cognition”.
- [458] Smart, P. R. Submitted(b). “Knowledge Machines”.
- [459] Smart, P. R., P. C. Engelbrecht, D. Braines, M. Strub, and C. Giammanco. 2010a. “The Network-Extended Mind”. In: *Network Science for Military Coalition Operations: Information Extraction and Interaction*. Ed. by D. Verma. Hershey, Pennsylvania, USA: IGI Global.
- [460] Smart, P. R., R. Heersmink, and R. W. Clowes. 2017. “The Cognitive Ecology of the Internet”. In: *Cognition Beyond the Brain: Computation, Interactivity and Human Artifice*. Ed. by S. J. Cowley and F. Vallée-Tourangeau. 2nd. Cham, Switzerland: Springer International Publishing.
- [461] Smart, P. R., T. D. Huynh, D. Braines, and N. R. Shadbolt. 2010b. “Dynamic Networks and Distributed Problem-Solving”. In: *Knowledge Systems for Coalition Operations (KSCO'10)*. Vancouver, British Columbia, Canada.

- [462] Smart, P. R., T. D. Huynh, D. Mott, K. Sycara, D. Braines, M. Strub, W. R. Sieck, and N. R. Shadbolt. 2009. "Towards an Understanding of Shared Understanding in Military Coalition Contexts". In: *3rd Annual Conference of the International Technology Alliance (ACITA '09)*. Maryland, USA.
- [463] Smart, P. R. and A. Madaan. in press. "The Social Scaffolding of Machine Intelligence". *International Journal On Advances in Intelligent Systems*.
- [464] Smart, P. R. and N. R. Shadbolt. 2014. "Social Machines". In: *Encyclopedia of Information Science and Technology*. Ed. by M. Khosrow-Pour. Hershey, Pennsylvania, USA: IGI Global.
- [465] Smart, P. R. and N. R. Shadbolt. in press. "The World Wide Web". In: *The Routledge Handbook of Applied Epistemology*. Ed. by J. Chase and D. Coady. New York, New York, USA: Routledge.
- [466] Smart, P. R., E. Simperl, and N. R. Shadbolt. 2014. "A Taxonomic Framework for Social Machines". In: *Social Collective Intelligence: Combining the Powers of Humans and Machines to Build a Smarter Society*. Ed. by D. Miorandi, V. Maltese, M. Rovatsos, A. Nijholt, and J. Stewart. Berlin, Germany: Springer.
- [467] Smart, P. R. and K. Sycara. 2013. "Collective Sensemaking and Military Coalitions". *Intelligent Systems*. 28(1): 50–56.
- [468] Smart, P. R. and K. Sycara. 2015. "Situating Cognition in the Virtual World". In: *6th International Conference on Applied Human Factors and Ergonomics*. Las Vegas, Nevada, USA.
- [469] Smith, M. N. 2016. "An iPhone Is an Extension of the Mind". http://www.slate.com/articles/technology/technology/2016/02/apple_and_the_fbi_think_iphones_are_safes_a_philosopher_explains_what_they.html. Accessed: 20th August 2016.
- [470] Someya, T., T. Sekitani, S. Iba, Y. Kato, H. Kawaguchi, and T. Sakurai. 2004. "A large-area, flexible pressure sensor matrix with organic field-effect transistors for artificial skin applications". *Proceedings of the National Academy of Sciences*. 101(27): 9966–9970.

- [471] Søraker, J. H. 2007. “The moral status of information and information technologies: A relational theory of moral status”. In: *Information Technology Ethics: Cultural Perspectives*. Ed. by S. Hongladarom and C. Ess. Hershey, Pennsylvania, USA: Idea Group Reference.
- [472] Sparrow, B., J. Liu, and D. M. Wegner. 2011. “Google Effects on Memory: Cognitive Consequences of Having Information at Our Fingertips”. *Science*. 333(6043): 776–778.
- [473] Spiers, H. J., E. Manley, R. Silva, T. Conroy Dalton, J. M. Wiener, C. Hölscher, V. Bohbot, and M. Hornberger. 2016. “Spatial navigation ability assessed in over 1 million people globally”. In: *Neuroscience*. San Diego, California, USA.
- [474] Stahl, G. and C. P. Rosé. 2011. “Group cognition in online teams”. In: *Theories of Team Cognition: Cross-Disciplinary Perspectives*. Ed. by E. Salas, S. M. Fiore, and M. P. Letsky. New York, New York, USA: Routledge.
- [475] Staley, D. J. 2014. *Brain, Mind and Internet: A Deep History and Future*. Basingstoke, England, UK: Palgrave Macmillan.
- [476] Stanney, K. M., R. S. Kennedy, J. M. Drexler, and D. L. Harm. 1999. “Motion sickness and proprioceptive aftereffects following virtual environment exposure”. *Applied Ergonomics*. 30(1): 27–38.
- [477] Stanney, K. M., D. D. Schmorrow, M. Johnston, S. Fuchs, D. Jones, K. S. Hale, A. Ahmad, and P. Young. 2009. “Augmented cognition: An overview”. *Reviews of Human Factors and Ergonomics*. 5(1): 195–224.
- [478] Steels, L. 2006. “Collaborative tagging as distributed cognition”. *Pragmatics & Cognition*. 14(2): 287–292.
- [479] Stepper, S. and F. Strack. 1993. “Proprioceptive determinants of emotional and nonemotional feelings”. *Journal of Personality and Social Psychology*. 64(2): 211–220.
- [480] Sterelny, K. 2003. *Thought in a Hostile World: The Evolution of Human Cognition*. Oxford, UK: Blackwell Publishing.

- [481] Sterelny, K. 2004. “Externalism, Epistemic Artefacts and the Extended Mind”. In: *The Externalist Challenge: New Studies on Cognition and Intentionality*. Ed. by R. Schantz. Berlin, Germany: de Gruyter.
- [482] Sterelny, K. 2007. “Social intelligence, human intelligence and niche construction”. *Philosophical Transactions of the Royal Society B: Biological Sciences*. 362(1480): 719–730.
- [483] Sterelny, K. 2010. “Minds: Extended or scaffolded?” *Phenomenology and the Cognitive Sciences*. 9(4): 465–481.
- [484] Strack, F., L. L. Martin, and S. Stepper. 1988. “Inhibiting and facilitating conditions of the human smile: A nonobtrusive test of the facial feedback hypothesis”. *Journal of Personality and Social Psychology*. 54(5): 768–777.
- [485] Strohmaier, M. and C. Wagner. 2014. “Computational social science for the World Wide Web”. *IEEE Intelligent Systems*. 29(5): 84–88.
- [486] Suarez, J. R. P. 2015. “We are Cyborgs: Developing a Theoretical Model for Understanding Criminal Behaviour on the Internet”. *Ph.D.* Huddersfield, West Yorkshire, UK: University of Huddersfield.
- [487] Sullins, J. 2016. “Information technology and moral values”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by E. N. Zalta. Spring 2016. Stanford, California, USA: Stanford University.
- [488] Sullivan, B. L., C. L. Wood, M. J. Iff, R. E. Bonney, D. Fink, and S. Kelling. 2009. “eBird: A citizen-based bird observation network in the biological sciences”. *Biological Conservation*. 142(10): 2282–2292.
- [489] Sumpter, D. J. T. and M. Beekman. 2003. “From nonlinearity to optimality: Pheromone trail foraging by ants”. *Animal Behaviour*. 66(2): 273–280.
- [490] Surowiecki, J. 2005. *The Wisdom of Crowds: Why the Many are Smarter than the Few*. New York, New York, USA: Random House.

- [491] Sutton, J. 2010. “Exograms and interdisciplinarity: History, the extended mind, and the civilizing process”. In: *The Extended Mind*. Ed. by R. Menary. Cambridge, Massachusetts, USA: MIT Press.
- [492] Swan, M. 2013. “The quantified self: Fundamental disruption in big data science and biological discovery”. *Big Data*. 1(2): 85–99.
- [493] Taivalsaari, A. and T. Mikkonen. 2017. “A Roadmap to the Programmable World: Software Challenges in the IoT Era”. *IEEE Software*. 34(1): 72–80.
- [494] Taraborelli, D. 2008. “How the Web is changing the way we trust”. In: *Current Issues in Computing and Philosophy*. Ed. by A. Briggie, K. Waelbers, and P. Brey. Amsterdam, The Netherlands: IOS Press.
- [495] Terzopoulos, D., X. Tu, and R. Grzeszczuk. 1994. “Artificial fishes: Autonomous locomotion, perception, behavior, and learning in a simulated physical world”. *Artificial Life*. 1(4): 327–351.
- [496] Thaler, S., E. Simperl, and K. Siorpaes. 2011. “SpotTheLink: A Game for Ontology Alignment”. *Wissensmanagement*. 182: 246–253.
- [497] Theiner, G. 2013. “Onwards and upwards with the extended mind: From individual to collective epistemic action”. In: *Developing Scaffolds in Evolution, Culture, and Cognition*. Ed. by L. R. Caporael, J. R. Griesemer, and W. C. Wimsatt. Cambridge, Massachusetts, USA: MIT Press.
- [498] Theiner, G. 2014. “Varieties of Group Cognition”. In: *The Routledge Handbook of Embodied Cognition*. Ed. by L. A. Shapiro. New York, New York, USA: Routledge.
- [499] Theiner, G., C. Allen, and R. L. Goldstone. 2010. “Recognizing group cognition”. *Cognitive Systems Research*. 11: 378–395.
- [500] Theiner, G. and T. O’Connor. 2010. “The Emergence of Group Cognition”. In: *Emergence in Science and Philosophy*. Ed. by A. Corradini and T. O’Connor. Abingdon, Oxon, UK: Routledge.

- [501] Thompson, C. 2013. *Smarter Than You Think: How Technology Is Changing Our Minds for the Better*. New York, New York, USA: The Penguin Press.
- [502] Thompson, E. and M. Stapleton. 2009. “Making sense of sense-making: Reflections on enactive and extended mind theories”. *Topoi*. 28(1): 23–30.
- [503] Thompson, J. B. 2011. “Shifting Boundaries of Public and Private Life”. *Theory, Culture & Society*. 28(4): 49–70.
- [504] Tiropanis, T., W. Hall, J. Crowcroft, N. Contractor, and L. Tassioulas. 2015. “Network Science, Web Science, and Internet Science”. *Communications of the ACM*. 58(8): 76–82.
- [505] Tollefsen, D. P. 2006. “From extended mind to collective mind”. *Cognitive Systems Research*. 7(2–3): 140–150.
- [506] Tomasello, M. 2000. “Primate cognition: Introduction to the issue”. *Cognitive Science*. 24(3): 351–361.
- [507] Tong, F. and M. S. Pratte. 2012. “Decoding patterns of human brain activity”. *Annual Review of Psychology*. 63: 483–509.
- [508] Tong, S. T., B. Van Der Heide, L. Langwell, and J. B. Walther. 2008. “Too much of a good thing? The relationship between number of friends and interpersonal impressions on Facebook”. *Journal of Computer-Mediated Communication*. 13(3): 531–549.
- [509] Trianni, V., E. Tuci, K. M. Passino, and J. A. R. Marshall. 2011. “Swarm cognition: An interdisciplinary approach to the study of self-organising biological collectives”. *Swarm Intelligence*. 5(1): 3–18.
- [510] Tribble, E. 2005. “Distributing Cognition in the Globe”. *Shakespeare Quarterly*. 56(2): 135–155.
- [511] Tribble, E. 2011. *Cognition in the Globe: Attention and Memory in Shakespeare’s Theatre*. New York, New York, USA: Palgrave Macmillan.
- [512] Tribble, E. and J. Sutton. 2011. “Cognitive ecology as a framework for Shakespearean studies”. *Shakespeare Studies*. 39: 94–103.

- [513] Tumasjan, A., T. O. Sprenger, P. G. Sandner, and I. M. Welpe. 2010. "Predicting elections with Twitter: What 140 characters reveal about political sentiment". In: *Fourth International AAAI Conference on Weblogs and Social Media*. Washington D.C., USA.
- [514] Turing, A. M. 1950. "Computing machinery and intelligence". *Mind*. 59(236): 433–460.
- [515] Turkle, S. 2011. *Alone Together: Why We Expect More From Technology And Less From Each Other*. New York, New York, USA: Basic Books.
- [516] Turner, J. S. 2011. "Termites as models of swarm cognition". *Swarm Intelligence*. 5(1): 19–43.
- [517] Tveit, Å. K. and A. Mangen. 2014. "A joker in the class: Teenage readers' attitudes and preferences to reading on different devices". *Library & Information Science Research*. 36(3): 179–184.
- [518] Valins, S. 1966. "Cognitive effects of false heart-rate feedback". *Journal of Personality and Social Psychology*. 4(4): 400–408.
- [519] Van Der Heide, B., J. D. D'Angelo, and E. M. Schumaker. 2012. "The effects of verbal versus photographic self-presentation on impression formation in Facebook". *Journal of Communication*. 62(1): 98–116.
- [520] Van Kleek, M. and K. O'Hara. 2014. "The Future of the Social is Personal: The Potential of the Personal Data Store". In: *Social Collective Intelligence: Combining the Powers of Humans and Machines to Build a Smarter Society*. Ed. by D. Miorandi, V. Maltese, M. Rovatsos, A. Nijholt, and J. Stewart. Berlin, Germany: Springer.
- [521] Verbeek, P.-P. 2009. "Ambient intelligence and persuasive technology: The blurring boundaries between human and technology". *Nanoethics*. 3(3): 231–242.
- [522] Verdoux, P. 2011. "Emerging technologies and the future of philosophy". *Metaphilosophy*. 42(5): 682–707.
- [523] Vierkant, T. 2014. "Mental muscles and the extended will". *Topoi*. 33(1): 57–65.

- [524] Vincent, J., L. Haddon, and L. Hamill. 2005. "The influence of mobile phone users on the design of 3G products and services". *The Journal of the Communications Network*. 4(4): 69–73.
- [525] Vygotsky, L. S. 1978. *Mind in Society: The Development of Higher Psychological Processes*. Cambridge, Massachusetts, USA: Harvard University Press.
- [526] Walther, J. B. 1992. "Interpersonal Effects in Computer-Mediated Interaction: A Relational Perspective". *Communication Research*. 19(1): 52–90.
- [527] Walther, J. B., B. Van Der Heide, L. M. Hamel, and H. C. Shulman. 2009. "Self-generated versus other-generated statements and impressions in computer-mediated communication: A test of warranting theory using Facebook". *Communication Research*. 36(2): 229–253.
- [528] Walther, J. B., B. Van Der Heide, S.-Y. Kim, D. Westerman, and S. T. Tong. 2008. "The role of friends' behavior on evaluations of individual's Facebook profiles: Are we known by the company we keep?" *Human Communication Research*. 34: 28–49.
- [529] Wang, Q. and C. Nass. 2005. "Less Visible and Wireless: Two experiments on the effects of microphone type on users' performance and perception". In: *SIGCHI Conference on Human Factors in Computing Systems*. Portland, Oregon.
- [530] Ward, A. F. 2013a. "One with the Cloud: Why People Mistake the Internet's Knowledge for Their Own". *Ph.D.* Cambridge, Massachusetts, USA: Harvard University.
- [531] Ward, A. F. 2013b. "Supernormal: How the Internet Is Changing Our Memories and Our Minds". *Psychological Inquiry*. 24(4): 341–348.
- [532] Warwick, K. 2015. "The Merging of Humans and Machines". In: *Neurotechnology, Electronics, and Informatics*. Ed. by A. R. Londeral, P. Encarnação, and J. L. P. Rovira. London, UK: Springer.

- [533] Waters, C. N., J. Zalasiewicz, C. Summerhayes, A. D. Barnosky, C. Poirier, A. Gałuszka, A. Cearreta, M. Edgeworth, E. C. Ellis, M. Ellis, C. Jeandel, R. Leinfelder, J. R. McNeill, D. Richter, W. Steffen, J. Syvitski, D. Vidas, M. Wapreisch, M. Williams, A. Zhisheng, J. Grinevald, E. Odada, N. Oreskes, and A. P. Wolfe. 2016. “The Anthropocene is functionally and stratigraphically distinct from the Holocene”. *Science*. 351(6269): pages.
- [534] Webb, B. 1996. “A cricket robot”. *Scientific American*. 275(6): 94–99.
- [535] Weber, W., J. M. Rabaey, and E. Aarts, eds. 2010. *Ambient Intelligence*. Berlin, Germany: Springer.
- [536] Wegner, D. M. 1987. “Transactive memory: A contemporary analysis of the group mind”. In: *Theories of Group Behavior*. Ed. by B. Mullen and G. R. Goethals. New York, New York, USA: Springer.
- [537] Wegner, D. M., R. Erber, and P. Raymond. 1991. “Transactive memory in close relationships”. *Journal of Personality and Social Psychology*. 61(6): 923–929.
- [538] Wegner, D. M. and A. F. Ward. 2013. “How Google is changing your brain”. *Scientific American*. 309(6): 58–61.
- [539] Weiser, M. 1991. “The computer for the 21st century”. *Scientific American*. 265(3): 94–104.
- [540] Wellman, B. 2012. “Is Dunbar’s number up?” *British Journal of Psychology*. 103(2): 174–176.
- [541] Westerman, D., P. R. Spence, and B. Van Der Heide. 2012. “A social network as information: The effect of system generated reports of connectedness on credibility on Twitter”. *Computers in Human Behavior*. 28: 199–206.
- [542] Wheeler, M. 2005. *Reconstructing the Cognitive World*. Cambridge, Massachusetts, USA: MIT Press.
- [543] Wheeler, M. 2011. “Thinking beyond the brain: Educating and building from the standpoint of extended cognition”. *Computational Culture*. 1.

- [544] Wheeler, M. 2013. “What Matters: Real Bodies and Virtual Worlds”. In: *SmartData: Privacy Meets Evolutionary Robotics*. Ed. by I. Harvey, A. Cavoukian, G. Tomko, D. Borrett, H. Kwan, and D. Hatzinakos. New York, New York, USA: Springer.
- [545] Whitaker, R. M., M. J. Chorley, and S. M. Allen. 2015. “New frontiers for crowdsourcing: The extended mind”. In: *48th Annual Hawaii International Conference on System Sciences*. Grand Hyatt, Kauai, Hawaii, USA.
- [546] Whiten, A. and R. W. Byrne, eds. 1997. *Machiavellian Intelligence II: Extensions and Evaluations*. Cambridge, UK: Cambridge University Press.
- [547] Wilson, R. A. 2000. “The mind beyond itself”. In: *Metarepresentations: A Multidisciplinary Perspective*. Ed. by D. Sperber. New York, New York, USA: Oxford University Press.
- [548] Wilson, R. A. 2004. *Boundaries of the Mind: The Individual in the Fragile Sciences: Cognition*. New York, New York, USA: Cambridge University Press.
- [549] Wilson, R. A. 2014. “Ten questions concerning extended cognition”. *Philosophical Psychology*. 27(1): 19–33.
- [550] Wilson, R. A. and A. Clark. 2009. “Situated Cognition: Letting Nature Take its Course”. In: *The Cambridge Handbook of Situated Cognition*. Ed. by P. Robbins and M. Aydede. Cambridge, UK: Cambridge University Press.
- [551] Wolf, M. and M. Barzillai. 2009. “The importance of deep reading”. *Educational Leadership*. 66(6): 32–37.
- [552] Wollmer, M. A., C. de Boer, N. Kalak, J. Beck, T. Götz, T. Schmidt, M. Hodzic, U. Bayer, T. Kollmann, and K. Kollwe. 2012. “Facing depression with botulinum toxin: A randomized controlled trial”. *Journal of Psychiatric Research*. 46(5): 574–581.
- [553] Wood, D., M. Zaidman, L. Ruth, and M. Hausenblas. 2014. *Linked Data: Structured Data on the Web*. Shelter Island, New York, USA: Manning Publications Co.

- [554] Woods, S. P., J. E. Iudicello, L. M. Moran, C. L. Carey, M. S. Dawson, and I. Grant. 2008. "HIV-associated prospective memory impairment increases risk of dependence in everyday functioning". *Neuropsychology*. 22(1): 110–117.
- [555] Xie, J., S. Sreenivasan, G. Korniss, W. Zhang, C. Lim, and B. K. Szymanski. 2011. "Social consensus through the influence of committed minorities". *Physical Review E*. 84(1).
- [556] Yokota, T., P. Zalar, M. Kaltenbrunner, H. Jinno, N. Matsuhisa, H. Kitanosako, Y. Tachibana, W. Yukita, M. Koizumi, and T. Someya. 2016. "Ultraflexible organic photonic skin". *Science Advances*. 2(4): e1501856.
- [557] Yu, L. L., J. V. Nickerson, and Y. Sakamoto. 2012. "Collective Creativity: Where we are and where we might go". In: *Collective Intelligence*. Cambridge, Massachusetts, USA.
- [558] Zagzebski, L. T. 1996. *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*. Cambridge, UK: Cambridge University Press.
- [559] Zakay, D. and R. A. Block. 1997. "Temporal cognition". *Current Directions in Psychological Science*. 6(1): 12–16.
- [560] Zalasiewicz, J., M. Williams, A. Haywood, and M. Ellis. 2011. "The Anthropocene: A new epoch of geological time?" *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*. 369(1938): 835–841.
- [561] Zeng, D., S. Guo, and Z. Cheng. 2011. "The Web of Things: A survey". *Journal of Communications*. 6(6): 424–438.
- [562] Ziemke, T. 2003. "What's that thing called embodiment". In: *25th Annual Meeting of the Cognitive Science Society*. Boston, Massachusetts, USA.
- [563] Zollman, K. J. S. 2010. "The epistemic benefit of transient diversity". *Erkenntnis*. 72(1): 17–35.