Minds and Machines manuscript No.
(will be inserted by the editor)

Situating Machine Intelligence within the Cognitive Ecology of the Internet

Paul R. Smart

Received: date / Accepted: date

Abstract The Internet is an increasingly important focus of attention for the philosophy of mind and cognitive science communities. This is partly because the Internet serves as an important part of the material environment in which a broad array of human cognitive and epistemic activities are situated. The Internet can thus be seen as an important part of the 'cognitive ecology' that helps to shape, support and (on occasion) realize aspects of human cognizing. Much of the previous philosophical work in this area has sought to analyze the cognitive significance of the Internet from a human perspective. There has, as such, been little effort to assess the cognitive significance of the Internet from the perspective of 'machine cognition'. This is unfortunate, because the Internet is likely to exert a significant influence on the shape of machine intelligence. The present paper attempts to evaluate the extent to which the Internet serves as a form of cognitive ecology for synthetic (machine-based) forms of intelligence. In particular, the phenomenon of Internet-situated machine intelligence is analyzed from the perspective of a number of approaches that are typically subsumed under the heading of situated cognition. These include extended, embedded, scaffolded, embodied and collective approaches to cognition. For each of these approaches, the Internet is shown to be of crucial relevance to machine cognition. Such insights help us to appreciate the role of the Internet in advancing the current state-of-the-art in machine intelligence.

Keywords Machine Intelligence · Machine Learning · Artificial Intelligence · Internet · Situated Cognition · World Wide Web

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

1 Introduction

Over the past two decades, the Internet has emerged as an important part of the material environment in which an ever-expanding array of human cognitive and epistemic activities are situated. This makes the Internet a system of considerable importance to the philosophy of mind and cognitive science. Fueled by developments in wireless technology and the emergence of a plethora of Internet-enabled devices (e.g., smartphones), the Internet has arguably become a critical element of our cognitively-relevant environment or 'cognitive ecology'. Many kinds of cognitive task, not to mention an increasing number of social activities, rely on the Internet at one point or another, and this makes the Internet a critical focus of attention for those who see environmental (i.e., extra-neural and extra-organismic) factors as playing an explanatorily-significant role in human cognizing (e.g., Clark 2008).

The view of the Internet as a form of cognitive ecology ties in nicely with a number of so-called ecological approaches to cognition (Bateson 1972; Malafouris 2013; Hutchins 2010; Tribble and Sutton 2011; Neisser 1997; Cooke et al 2004; Hirose 2002; Barrett 2011). These approaches emphasize the role of agent-environment interactions in shaping, supporting, and perhaps even realizing human cognitive states and processes. Hutchins (2010), for example, embraces an ecological approach to cognition when he suggests that our attempt to understand "cognitive phenomena must include a consideration of the environments in which cognitive processes develop and operate" (p. 706). In the absence of this sort of wider 'ecological' analysis, Hutchins fears that we may lose sight the forces and factors that undergird our cognitive capabilities. In this respect, Hutchins expresses a view that is broadly consistent with approaches that are typically subsumed under the general heading of situated cognition (Robbins and Aydede 2009a).

The idea that the Internet forms part of the cognitive ecology for human agents is apparent in a number of works that are spread across the disciplines of philosophy, computer science and cognitive science. The idea has perhaps its most explicit expression in a paper by Smart et al. (in press). Here, the appeal to a cognitive ecological conceptualization of the Internet is used to motivate a situated approach to the analysis of human-Internet interactions. In particular, Smart et al. (in press) seek to assess the cognitive significance of the Internet by viewing human-Internet interactions from the perspective of extended (Clark 2008; Clark and Chalmers 1998), embedded (Rupert 2004, 2009), embodied (Shapiro 2007, 2011, 2014; Anderson 2003), and collective/distributed (Theiner 2014; Theiner et al 2010; Hutchins 2001, 1995a) approaches to cognition. Similar sorts of analyses are provided by a growing body of work in philosophy and cognitive science (Staley 2014; Halpin 2013; Halpin et al 2010; Smart 2012, 2013, 2014, in press; Clowes 2015).

Although the Internet is clearly an increasingly popular target of philosophical theorizing, there is one respect in which the existing literature is somewhat lacking. This shortcoming relates to the almost exclusive focus on human cognitive capabilities. This focus is, of course, understandable given the nature

of our interaction with (and increasing dependence upon) the contemporary Internet. The focus is, however, unfortunate, because their is considerable evidence to suggest that the Internet and Web are exerting profound effects on what might be broadly construed as machine intelligence or machine cognition. We can get a feel for the nature of this influence by reflecting on the way in which the advent of the Internet has influenced research into machine learning and artificial intelligence. We thus see increasing amounts of research being invested into cognitive computing systems (Hurwitz et al 2015; Kelly and Hamm 2013), deep learning algorithms (Najafabadi et al 2015), big data analytics (Hurwitz et al 2015; O'Leary 2013), the Semantic Web (Berners-Lee et al 2001), and so on. The Internet has also spawned a renewed interest in the familiar idea of human-machine symbiosis (e.g., Licklider 1960; Jacucci et al 2014), where human and machine capabilities are factored into hybrid processing loops that span the biological and technological domains. One finds evidence of this interest in the recent explosion of research into 'human-in-the-loop' systems (Branson et al 2014), game-powered machine learning systems (Barrington et al 2012), human computation systems (Law and von Ahn 2011; Michelucci 2013), citizen science platforms (Khatib et al 2011a) and Internet-enabled forms of collective intelligence (Michelucci and Dickinson 2016). This efflorescence of research relating to Internet-enabled forms of machine intelligence is largely overlooked by existing attempts to assess the cognitive impacts of the Internet. And yet, when it comes to cognitive capabilities, it seems that the Internet is likely to be just as relevant to machine cognition as it is to human cognition (perhaps even more so).

In view of all this, in the present paper, I want to suggest that we should see the Internet as an environment that is relevant to the cognitive capabilities of *both* human and machine agents. Such a view encourages us to assess the cognitive impacts of the Internet from the conceptual standpoint of situated cognition (Robbins and Aydede 2009b)[1]. In what follows, therefore, I will seek to explore the way in which the notions of extended, embedded, embodied, scaffolded, and collective cognition can be used to help us gain a better understanding of online or Internet-situated forms of machine intelligence[2]. Not only does this approach help to establish an important and interesting parallel with work that focuses almost exclusively on human agents, it also helps, I suggest, to reveal a number of important directions for future research into machine intelligence.

[1] This is because a situated approach to cognition enables us to focus on the role of the extra-organismic and extra-neural environment (i.e., the features of the cognitive ecology) in shaping the cognitive capabilities of the biological or (in this case) the technological agent.

[2] It should be noted that this review is not exhaustive, in the sense that it does not attempt to consider a number of other forms of cognition that might be subsumed under the banner of situated cognition. There is, in particular, no attempt to cover the considerable body of work relating to enactive cognition (Stewart et al 2010; Engel et al 2016; Froese et al 2011) and grounded cognition (Barsalou 2008, 2010; Pezzulo et al 2013; Pecher and Zwaan 2005).

2 Machine Cognition?

In attempting to formulate a situated account of machine cognition, it is obviously important to explain what is meant by the notion of 'machine cognition'. In the human case, it seems natural to talk of our information processing capabilities in cognitive terms. Such, however, is not the case (or at least not always the case) when it comes to instances of machine-based information processing. Should the performance of a computer vision system or a machine reasoning system be glossed in cognitive terms? If not, then the present attempt to apply the conceptual apparatus of situated cognition to the case of machine intelligence may be something of a non-starter. Ideally, what is required here is an understanding of what Adams and Aizawa (2010) refer to as the 'mark of the cognitive': the set of features that enable us to discriminate the cognitive from the non-cognitive. Only in the light of that understanding, perhaps, will we be able to determine if some instance of machine-based processing should be graced with the label 'machine cognition'.

Unfortunately, there is (as yet) no definitive answer to the question of what it is that makes a particular process a cognitive process. As Clark (2010a) notes, perhaps the best that can be done in this situation is a take a 'cognition is as cognition does' approach—seeing the realm of the cognitive as tied to the way in which particular processes help to support the expression of intelligent behaviour. We should thus, perhaps, recognize an instance of cognition based on the characteristic *effects* of a particular process rather than its characteristic *causes* (see Clark 2010a). A somewhat similar view is proposed by Rowlands (2006). He suggests that a cognitive process should be seen as a process that is "(i) required for the accomplishing of a cognitive task, (ii) involves information processing, and (iii) is of the sort that is capable of yielding a cognitive state" (p. 32). However, as is perhaps all too obvious, there are number of problems with this definition. Firstly, it introduces two new terms (i.e., 'cognitive task' and 'cognitive state'), neither of which are particularly well-understood. In addition, Rowlands (2006) suggests that we can resort to ostensive definitions in individuating cognitive tasks from their non-cognitive counterparts. However, inasmuch as we accept that cognitive tasks can be afforded ostensive definitions, then it is unclear why a similar (ostensive) approach could not be used to assist with the distinction between cognitive and non-cognitive processes.

My own view, here, is that we should to regard our ability to individuate the realm of the cognitive by linking it to the way in which certain types of processes contribute to the expression of certain kinds of behaviour. I thus concur with Clark (2010a) when he notes that the notion of a cognitive process "might be best unpacked as the notion of a process that supports certain kinds of behaviour (actual and counterfactual)" (p. 93). This sort of approach to cognition is, I suggest, one that is perfectly compatible with the possibility of machine cognition. Indeed the current interest in cognitive computing systems (Hurwitz et al 2015; Kelly and Hamm 2013; Modha et al 2011) is testament to the fact that at least *some* forms of machine-based processing can be (and

typically are) cast in cognitive terms. All that the notion of machine cognition requires, it seems, is that we encounter a process that we are prepared to view as cognitive in nature. A suitable test here—and one that helps to avoid any lingering form of biological bias—is to adopt a form of parity principle for machine cognition[3]:

> **Parity Principle for Machine Cognition**
> Whenever we encounter a putative case of machine cognition, we should ask ourselves whether the target process is one that we would be content to label as a cognitive process if such a process were—contrary to all the facts—to have taken place within the head of an individual human agent.

What the parity principle asks us to do, in other words, is to imagine a state-of-affairs in which some episode of machine-based processing (e.g., machine vision) is subject to neurological realization. With this state-of-affairs in mind, we can then ask ourselves whether we would be happy to see this imaginary case of information processing as a *bona fide* cognitive process. If the answer to this question is 'yes', then why should we demur from the conclusion that the same process, as realized by some alternative nexus of material elements, is not a cognitive process?

3 Extended Cognition

The term 'extended cognition' identifies an important and influential body of work within the philosophy of mind (see Clark 2008). At its core is the claim that the causally-active physical vehicles of cognition can, at least in principle, extend beyond the biological borders of the individual and include a range of extra-organismic (perhaps even non-biological) resources (Clark 2008; Clark and Chalmers 1998).

Although the notion of extended cognition is intended to challenge our bio- and neuro-centric intuitions (and prejudices) regarding the material bases of intelligence, there is a sense in which the existing philosophical debate remains skewed towards the biological realm. We see this in the way in which practically all discussions of extended cognition limit their attention to the case of human (or at least biological) intelligence. In the majority of cases, therefore, what emerges as the central focus of philosophical debate is whether *human* cognition is materially extended as a result of some form of bio-artifactual or bio-technological bonding. This focus is, of course, understandable, given the inherent interest in explaining the rather distinctive forms of cognitive success that are the hallmark of our species. It is, however, a focus that is challenged by the increasing sophistication of machine learning systems and cognitive computing platforms. And it is further challenged, I would argue, once we situate machine intelligence within the cognitive ecology of the Internet.

[3] This is an adaptation of the parity principle that is used by Clark and Chalmers (1998) to motivate the case for extended cognition.

The idea that I want to canvass, here, is called Human-Extended Machine Cognition (HEMC). As its name suggests, HEMC is a form of extended cognition. However, rather than being centered on an individual human agent, it is instead centered on some form of synthetic intelligent system. HEMC is, in essence, the idea of materially-extended machine cognition, where the relevant information processing loops extend beyond the reams of machine-based computational processing to include the contributions of one or more human agents. Whereas conventional cases of extended cognition oblige us to focus on the individual human agent and consider how their cognitive states and processes might be extended by virtue of their interactions with a surrounding nimbus of non-biological elements, the HEMC concept encourages us to engage in something of a 'mental flip', wherein we focus our attention on the way in which the cognitive processes of some intelligent machine-based system (e.g., a cognitive computing platform) is realized, at least in part, by a more-or-less permanent penumbra of human individuals.

Inasmuch as we accept the basic possibility of HEMC, we should obviously consider the extent to which such forms of cognitive extension are materially possible. Although it should be clear that nothing prevents important and compelling cases of HEMC arising outside the socio-technical context of the Internet[4], I want to suggest that the Internet affords an environment that is particularly suited to the practical realization of HEMC-based systems. One reason for this concerns the growth and popularity of the Internet. Crucially, the Internet has emerged as a major part of the technological fabric associated with contemporary society. It is, as such, a system that is used by an increasing proportion of the World's population. The result is that we can see the Internet as providing access to a large number of cognitively-sophisticated processing elements (i.e., individual human agents) that are more-or-less continuously available[5]. Although we might question the extent to which the elements of this large-scale social environment are suitably poised to participate in episodes of machine-based processing[6], what should not be in doubt is the fact that

[4] For example, we can imagine human biological brains beng incorporated into complex computational tasks by virtue of the use of advanced brain activity recording devices or imaging techniques. The use of such techniques in the context of rapid serial visual presentation tasks provides one example of such forms of bio-technologically hybrid processing (see Huang et al 2011).

[5] See Lasecki and Bigham (2013) for an overview of issues relating to the real-time and continuous exploitation of the human social environment.

[6] There is, of course, an important issue hereabouts concerning the extent to which humans agents are actually willing to participate in machine-based processing. The mere fact that almost anyone on the planet can connect to the Web does not mean that anyone will actually be 'available' for incorporation into an information processing routine. Issues of incentivization and motivation are clearly important here, as is the extent to which the temporal profile of human participation can be aligned with the temporal constraints of the target task. With respect to issues to timing, it is important to note that the real-time recruitment and engagement of human participants into some form of information processing activity is a topic that is the focus of current research efforts in the human computation and crowdsourcing communities. Research in this area typically goes under the heading of real-time or continuous crowdsourcing (Lasecki and Bigham 2013).

the Internet provides the basis for large-scale forms of social participation in a variety of online tasks (see Kraut et al 2010). Indeed, existing work in the area of human computation (Law and von Ahn 2011; Michelucci 2013) already attests to the possibility of 'incorporating' human individuals into bio-technologically hybrid processing loops.

There are a number of reasons why we should take the notion of HEMC seriously. Firstly, the notion serves as a conceptual bridge between two currently disparate areas of intellectual inquiry. One of these areas is, of course, delineated by philosophical debates concerning extended cognition and the extended mind. The other area emerges in relation to research on Internet-enabled forms of human-machine cooperation (Smart and Shadbolt 2014; Michelucci 2016). This includes, for example, a considerable body of work that has emerged in relation to human computation systems (Law and von Ahn 2011; Michelucci 2013) and citizen science platforms (Lintott and Reed 2013). Indeed, when reading the literature on human computation, one sometimes encounters an almost implicit reference to the idea of multiple (human) individuals serving as part of the realization base for the processing capabilities of socio-computational systems (see Michelucci 2016). The notion of HEMC provides a means of making this idea explicit, while simultaneously establishing an important point of conceptual contact with recent work in the philosophy of mind (Clark 2008).

Another reason why the notion of HEMC is important is because of its potential to illuminate novel approaches to the implementation of intelligent systems. Perhaps one of the most important considerations here relates to issues of complementarity. It has long been recognized that one of the virtues of cognitive extension is its ability to harness the complementary contributions of both biological and non-biological resources (Sutton 2010; Wilson and Clark 2009). Wilson and Clark (2009), for example, note the importance of extended cognition in yielding "hybrid processes in which the inner and the outer contributions are typically highly distinct in nature, yet deeply integrated and complementary" (p. 72). An appreciation of such complementarity is, of course, also evident in much of the literature pertaining to machine intelligence, especially work that goes under the heading of complementary computing (Kapoor et al 2008), mixed-initiative systems (Horvitz 2007), interactive machine learning (Fogarty et al 2008) and heterotic computing (Kendon et al 2015). There is, however, a sense in which the notion of HEMC adds something new to this debate on complementarity. This is the idea—often encountered in the extended mind literature—that the exploitation of complementary, bio-external resources may be of crucial relevance in accounting for some of the more advanced aspects of human-level cognition. As Clark (2003) notes "what is special about human brains, and what best explains the distinctive features of human intelligence, is precisely their ability to enter into deep and complex relationships with non-biological, props, and aids" (p. 5). Such claims should encourage the designers of intelligent systems to give serious thought to issues of hybridization and complementarity. For inasmuch as human cognition is grounded in the ability to assemble and then exploit processing loops that

extend across the brain, the body and the world, then perhaps it is somewhat unfair to expect a system that trades in only one kind of information processing economy (e.g., a connectionist or classical rule-and-symbol processing regime) to achieve the kinds of cognitive success that we deem to be representative of human-level intelligence. Put another way, if hybridization is a fundamental feature of the human cognitive architecture, then perhaps our best approach to engineering intelligent systems is to rely on solutions that factor in the representational and computational contributions of elements that are drawn from *both* the biological and non-biological domains. Just as advanced forms of human intelligence are mostly (and possibly always) grounded in networks of bio-artifactual and bio-technological coupling, so too perhaps the route to machine intelligence lies in an ability to merge the modes of representation and processing made available by both our biological brains and our current arsenal of non-biological computing systems.

4 Embedded Cognition

According to those who embrace the notion of extended cognition, human cognitive processes sometimes supervene on elements that lie outwith the biological borders of the individual. This is clearly a somewhat radical-sounding claim. A more conservative view comes in the form of embedded cognition, which is championed by the likes of Robert Rupert (2004; 2009). In contrast to the idea that extra-organismic elements play a constitutive role in the realization of cognitive states and processes, Rupert suggests that we should see bio-external resources as merely influencing inner, biologically-based(?) cognitive routines. This is not to say that advocates of embedded cognition do not recognize the importance of extra-organismic resources in shaping the course of cognitive processing. As Rupert (2004) himself notes:

> "...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" (p. 393).

As this quotation makes clear, proponents of embedded cognition are keen to emphasize the way in which cognitive states and processes depend on aspects of the extra-organismic environment, even though they reject the claim that such resources should be seen as part of the physical fabric that realizes such states and processes.

This is clearly not the place to revisit arguments concerning the relative merits of extended vs. embedded conceptions of cognition (see Clark 2008, 2007, 2010b). What is important, in the present context, is the extent to which the notion of embedded cognition helps to improve our understanding of the space of possibilities regarding the design and implementation of Internet-situated forms of machine intelligence. In fact, as is noted by Clark (2008, p. 139) we should not see the choice between embedded and extended cognition as a "zero-sum game". Instead, we should see the embedded and extended

perspectives as conceptual lenses that are "apt to draw attention to certain features, regularities, and contributions while making it harder to spot others or to give them their problem-solving due" (Clark 2008, p. 139). With this in mind, in what follows, I will attempt to highlight a couple of areas where I believe the embedded approach is to be preferred over the extended alternative, specifically when it comes to the notion of Internet-situated machine intelligence.

One situation in which the human social environment plays an important role in shaping the course of machine-based processing is to be found in the context of collaborative filtering systems. Consider, for example, the approaches adopted by companies such as Amazon and Netflix with regard to their product recommendation systems. In both cases, there is a significant problem to be solved, namely, the identification of those resources that might appeal to users based on their history of purchasing behaviour. One approach to this problem is to engage in a detailed analysis of each resource, attempting to judge the relative similarity of each resource with respect to a number of dimensions of interest. This, it should be clear, is a task that is of sufficient scale and complexity as to challenge the current state-of-the-art in machine intelligence[7]. The alternative approach is, of course, the one embraced by both Amazon (Linden et al 2003) and Netflix (Amatriain 2012): simply monitor the purchasing behaviour of the user community and rely on the use of collaborative filtering techniques to make future product recommendations.

What we encounter in the case of collaborative filtering, I suggest, is a situation in which particular kinds of machine intelligence lean very heavily on the human social environment. However, in this particular case, it does not seem appropriate to regard the human social environment as a constitutive part of the physical fabric that realizes the processes associated with the target task (i.e., the derivation of product recommendations). Instead, it seems far more appropriate to see the human social environment as (e.g.) generating representations that productively re-shape and re-configure the nature of the computational problem that needs to be tackled by the technological system.

Examples of this sort of (socially-)embedded machine intelligence are, in fact, relatively commonplace on the Internet. Google's PageRank algorithm (Brin and Page 1998), for example, which sits at the heart of its market-leading search capability, is grounded in the linking behaviour of the Web user community: as links accrue to particular resources, they provide an indication of the worth, interest or value of those resources with respect to subsequent information retrieval activities. In the absence of this sort of structural configuration and enrichment of the online environment, the PageRank algorithm would not be able to support the ranking and sorting capabilities that are manifested by the search engine. Similarly, when it comes to judgments about the reliability of particular bodies of online information, the endorsement ac-

[7] To see this, we need only think of the diversity of products offered by a site like Amazon. Similarly, when it comes to Netflix, it should be clear that the automatic classification and ranking of movie titles according to a range of abstract semantic categories and dimensions (e.g., 'romantic comedy') is a task that is far from straightforward.

tivities of human users (in the form of, for example, explicit user ratings) provides an important source of information that can be factored into credibility evaluations (Taraborelli 2008).

The moral to emerge from these sorts of cases is that the human social environment can often play an important (perhaps indispensable) role in enabling some form of (synthetic) intelligent processing capability to emerge. Unlike the case of HEMC, the human agents in these cases do not seem to be actively involved in the realization of task-relevant information processing routines. Instead, they seem to engage in a productive re-configuration or enrichment of the environment in which machine-based processing takes place. The image that emerges, here, is somewhat akin to that proposed by David Kirsh (1995). Kirsh suggests that human agents engage in the active structuring of their environments in order to simplify or transform the nature of the problem that confronts the biological brain. Something similar can perhaps be said to occur in the aforementioned cases of product recommendation, resource ranking and reliability evaluation. In such cases, machine-based capabilities can be said to depend on the social environment as a means of simplifying or transforming the kind of problem that needs to be solved by the technological system. The nature of this dependency, it should be clear, is significant. Absent the human contributions and the performance of the machine-based system is likely to suffer as a result. Indeed, in many cases, the nature of the computational task that confronts the machine-based system is one that can only be achieved with the sort of support (deliberate or otherwise) that is provided by the human user community.

5 Scaffolded Cognition

Scaffolded cognition, as it name suggests, is a form of cognition in which an array of extra-organismic elements are seen to play an important role in the development of certain kinds of cognitive capability. The term has its roots in the developmental psychology literature, most notably the work of the Soviet psychologist, Lev Vygotsky (1986). Vygotsky's contribution was to recognize the role played by extra-organismic resources, especially the human social environment, in shaping the developmental profile of a young human infant. The core idea was that the provision of support at crucial points in the developmental trajectory of a young child would eventually enable the child to achieve success in the absence of such support. The support provided by an adult caregiver while a young infant takes her first few faltering steps is one example of this sort of scaffolded development.

There is, of course, a certain sense in which scaffolded cognition appears very similar to both extended and embedded cognition. As with the notions of extended and embedded cognition, scaffolded cognition is a form of cognition that is seen to rely on a surrounding nexus of extra-organismic elements. What distinguishes scaffolded cognition from both extended and embedded cognition, however, is the role played by the extra-organismic elements in the

development of cognitive capabilities. In the case of scaffolded cognition, a supportive framework is erected around a developing structure or capability in order to support its full emergence. Once the capability is fully developed, the scaffolding is no longer required and can be dispensed with. This, I suggest, marks an important point of difference between the notions of scaffolded cognition on the one hand and the notions of extended and embedded cognition on the other[8]. In cases of scaffolded cognition, the environmental resources are helping to alter and inform the development of a set of biologically-based cognitive processing routines—routines that will, at some later stage, emerge as free-standing. This is not something that the proponent of either extended or embedded cognition needs to commit to. In the case of extended cognition, for example, the function of bio-external resources is not to scaffold the development of a biologically-based processing ability (although that may very well be one of their side-effects); neither is it to help bring about a future state-of-affairs in which the bio-external resources are no longer required for the runtime realization of some form of cognitive competence. Rather, the purpose of cognitively-potent bio-technological mergers is often to give birth to a new kind of capability, one that is perhaps forever beyond the developmental reach of the biologically-based cognizer.

Analogues to the notion of scaffolded cognition can, I suggest, be found in the literature pertaining to Internet-situated forms of machine intelligence. This is not, in fact, as difficult as it might appear, since we have already seen how the Internet provides the basis for ever-more intimate forms of cognitive contact with the human social environment (see Section 3), and it is typically elements of the human social environment (i.e., individual human agents) that provide the crucial forms of support in the case of conventional (i.e., human-centered) forms of scaffolded development (recall the role of human caregivers in supporting the infant's early ambulatory endeavors)[9]. What we are after in the case of machine-based scaffolded cognition, therefore, is a state-of-affairs in which human agents are providing a form of temporary support to the development of some 'mature' machine-based capability. An example of this can be found, I suggest, in the literature on citizen science systems (Lintott and Reed 2013). Consider, for example, a citizen science system known as Galaxy Zoo. Galaxy Zoo is a system in which human participants are asked to identify and classify galaxies based on images from the Sloan Digital Sky Survey (Lintott et al 2008). The basis for human involvement, here, relates to the difficulty of identifying the taxonomically-relevant features of galaxies from a set of 2D images. Recognizing that a particular galaxy is a spiral galaxy, for example,

[8] As an aside, I suggest that this difference serves to undermine claims that scaffolded cognition can be used a conceptual substitute for extended cognition (see Sterelny 2010). If the concept of scaffolded cognition represents a distinct form of cognition, then it is difficult to see how it is able to accommodate the claims advanced by the proponents of extended cognition.

[9] This should not, of course, blind us to the possibility that machine intelligence might be scaffolded in ways that do *not* involve human agents. There might, for example, be cases where the development of machine-based cognitive capabilities are scaffolded by other kinds of technological system.

requires an ability to detect an array of morphological features from a complex (and highly variable) set of visual cues. Humans, as it turns out, are quite good at this task, but the task presents something of a challenge for machine-based vision systems. Progress in machine vision tasks often comes about as the result of the application of machine learning techniques to large bodies of appropriately labeled trailing data. In the case of many machine vision tasks, however, the relevant body of training data is simply not available. The task of galaxy classification is no exception, here. But now notice something important: by virtue of the classification efforts of human volunteers, the raw data for the galaxy classification task (i.e., the set of 2D images) is altered. In place of a set of unlabeled images, we now have a growing corpus of images that are reliably associated with taxonomically-relevant labels. The availability of this *enriched* dataset opens up a range of, previously unavailable, opportunities for machine-based processing. One such opportunity is, of course, to exploit the annotated image set in the context of machine learning techniques that aim to enhance the performance of automated image classifiers. This is precisely the strategy adopted by Dieleman et al. (2015). Dielemen et al. sought to apply a novel class of neural networks, called deep convolutional neural networks, to the galaxy classification task. Here, the availability of the training dataset, originating from the actions of Galaxy Zoo volunteers, was instrumental in supporting the emergence of a novel form of machine-based intelligent processing capability—one which we can (and should) recognize as a genuine case of machine cognition, and one which (in my view) also represents a *bona fide* case of scaffolded cognitive development[10].

The importance of human contributions to the development of machine learning resources is explicitly recognized by those who seek to engage human subjects in computationally-difficult tasks. With respect to citizen science systems, for example, Lintott and Reed (2013) note that one of the limiting factors in the development of automated processing solutions is the availability of sufficiently well-structured training datasets, and that one of the key advantages of citizen science projects is the provision of such datasets. Similarly, when it comes to a class of systems known as Games With A Purpose (GWAP)[11], von

[10] It should be clear, for example, that the human contributions are playing an important (and perhaps indispensable) role in supporting the emergence of what is ultimately a free-standing capability. There is, in addition, no reason why we should fail to recognize this as a genuine case of machine cognition. Invoking the test outlined in Section 2, we can ask ourselves whether the machine-based capability is one that we would recognize as cognitive if it were to be performed by a human agent. The result, in this case, is surely positive. After all, the task that is being performed by the machine vision system—i.e., morphological galaxy classification—is more-or-less exactly the same task as that being performed by the human agents who generated the training data. If we accept that classification outcomes in the human case are the result of the exercise of a cognitive ability, why should we regard the outcomes of machine-based processing any differently?

[11] GWAPs are computer games that are designed to make use of game player actions (von Ahn 2006; Savage 2012; Cooper et al 2010b). Often such games are used to support the process of scientific discovery (Good and Su 2011; Khatib et al 2011b). We thus encounter examples of GWAPs in the domains of neuroscience (Marx 2013), proteomics (Cooper et al 2010a), astronomy (Lintott et al 2008), and oncology (Coburn 2014).

Ahn and Dabbish (2008) are keen to stress the role of human contributions in giving rise to ever-more intelligent forms of machine-based processing:

> "By leveraging the human time spent playing games online, GWAP game developers are able to capture large sets of training data that express uniquely human perceptual capabilities. This data can contribute to the goal of developing computer programs and automated systems with advanced perceptual or intelligence skills" (von Ahn and Dabbish 2008, p. 67).

In some cases, the machine-based elements can even play an active role in shaping the course of their own cognitive development. They do this by soliciting particular kinds of input from the human community. A good example of this kind of 'self-structured machine learning' or 'active learning' (see Settles 2012) is a system described by Barrington et al. (2012). The system in question uses a socially-oriented online game, called Herd It, in which groups of human individuals annotate a musical resource with descriptive tags. These annotations are subsequently used to train a supervised machine learning system that ultimately aims to perform the annotation task independently of the human agents. All this, of course, is broadly consistent with the general profile of machine learning. But what makes Barrington et al.'s system of particular interest is the way in which the machine *directs* the course of its own learning. It does this by actively selecting the musical resources that will be the focus of tagging efforts by the human game players. This is important because it gives the machine an opportunity to select those forms of feedback that are likely to be of greatest value in terms of informing its subsequent cognitive development. In the words of Barrington et al. (2012), "the machine learning system actively directs the annotation games to collect new data that will most benefit future model iterations" (p. 6411).

The cases of Galaxy Zoo and Herd It are thus compelling examples of how human contributions can play an important, and perhaps indispensable, role in the development of machine-based capabilities. Such cases should, I suggest, be viewed as the machine-based counterparts of the more conventional cases of scaffolded cognition that are discussed in the developmental psychology (Vygotsky 1986) and cognitive science (Sterelny 2010) literatures.

The scaffolding of machine-based capabilities helps to highlight the transformative potential of the Internet with respect to the realization of new kinds of cognitive processing capabilities. In this respect, the conceptual value of scaffolded cognition is similar to that identified in the case of extended cognition. In both cases we are able to see how the Internet enables the human social environment to extend the reach of machine-based cognitive capabilities. In the case of extended cognition, the focus, of course, is on the extent to which human agents can be incorporated into the cognitively-relevant processing routines of a machine-based system. In view of the current discussion of scaffolded cognition, however, it should be clear that the transformative impact of the Internet is not exhausted by these sorts of intimate symbiotic cognitive unions. We can also now appreciate the role of online social participation in shaping

the developmental trajectory of (what are ultimately) free-standing machine-based capabilities. Once we recognize this possibility, it alters the way in which we approach the *development* of online intelligent systems. Thus, rather than simply designing systems with the aim of exploiting human capabilities for the purposes of performing an otherwise intractable task, we can *also* think about the way in which certain kinds of human inputs may function as a form of cognitive scaffolding, working to enhance machine-based capabilities to the point at which human input is no longer required. This, I suggest, is one of the virtues of viewing the Internet through the conceptual lens of scaffolded cognition: it enables us to appreciate the sort of opportunities the Internet provides for the development of machine intelligence and adjust our development approaches accordingly[12].

In addition to all this, however, the notion of scaffolded cognition enables us to establish what I believe is a productive form of conceptual contact with work in cognitive developmental science. One point of contact comes in respect of work that shows how maturational shifts in cognitive, sensory and motor capabilities may be of crucial relevance to the emergence of advanced forms of cognitive competence (Gómez et al 2004; Newport 1990; Elman 1993; Nagai et al 2011; Bjorklund 1997). Such ideas are typically encountered in the context of human language learning. Bjorklund (1997), for example, suggests that by imposing a set of constraints on the kind of information structures that can be processed, maturational mechanisms can be seen as supporting the progressive reshaping of the 'effective' structure of a human infant's linguistic environment, transforming what might seem like an impossible language learning task into something a little more congenial. Similar ideas can be found in more recent work in developmental robotics. Gómez et al. (2004), for example, describe an intriguing set of results concerning the development of sensorimotor capabilities in a real-world robotic system. They report that a developmental profile characterized by progressive increments in the complexity of sensory, motor and neurocomputational subsystems results in a profile of task performance that is superior to that of a robot in which the relevant maturational processes are disabled. Commenting on this developmentally-grounded dissociation in 'adult' performance profiles, they suggest that:

> "...rather than being a problem, early morphological and cognitive limitations effectively decrease the amount of information that infants have to deal with, and may lead to an increase in the overall adaptivity of the organism" (Gómez et al 2004, p. 119).

Transposing all this to the realm of Internet-situated machine intelligence, we encounter a potentially important research opportunity. This relates to the

[12] One example where we see this sort of influence is in respect of the development of citizen science systems. In the case of PenguinWatch (see https://www.penguinwatch.org/), for example, human volunteers are required to identify the location of penguins in a set of images using a distinctive annotation. Such annotations provide a form of additional structure that can be used to bolster the performance of machine learning algorithms (Grant Miller, pers. comm.).

extent to which maturational shifts in computational parameters can work in concert with the scaffolding provided by the human social environment in order to enhance the capabilities of at least certain kinds of intelligent system[13]. There are clearly a number of ways in which this sort of dynamic dovetailing of intrinsic capabilities with environmental scaffolding might be realized. One possibility is for a machine-based learning system to exert some degree of control over the sort of scaffolding that is supplied by the human social environment, in the manner, perhaps, of the active learning system described by Barrington et al. (2012)[14]. Alternatively, a system could be configured so as to process inputs of increasing complexity as learning progresses[15]. What is important, here, is not so much the details of how the Internet could be used to scaffold the development of machine intelligence, so much as the idea that a number of new lines of scientific enquiry come into clearer view once we situate the notion of scaffolded cognition at the heart of our conceptual vista. By viewing the Internet as a form of cognitive ecology, and by looking at machine intelligence through the conceptual lens of scaffolded cognition, we are able to identify a set of relatively new routes by which advances in machine intelligence could be accomplished. In particular, what the notion of scaffolded cognition gives us is an alternative way of viewing the cognitive significance of human contributions to the online environment. Rather than always requiring the human community to engage in tasks that are beyond the ken of machine-based systems, we can begin to see the human community as playing a more indirect role in the creation of resources and structures that help to shape the developmental trajectory of machine-based cognitive abilities. In parallel with the notions of scaffolded development in the human case, it is possible that such forms of scaffolding are of crucial relevance in enabling machines to establish a firmer grip on cognitive capabilities that might otherwise have been forever beyond their grasp.

6 Embodied Cognition

Issues of material embodiment and environmental embedding typically find their greatest expression in an important and influential body of research that goes under the heading of 'embodied cognition'. Although there are a

[13] Obviously, not all intelligent systems are poised to benefit from a combination of maturational changes and scaffolded development. It should be relatively clear that only systems exhibiting some degree of 'plasticity' (with respect to their configuration) are able to benefit from scaffolded development.

[14] Active learning systems can thus be seen to implement a form of 'adaptive' (Smart et al 2010) or 'autonomous coupling' (Clark 1997, pp. 133-135). In this case, an agent has some degree of 'metacognitive' control over the time-variant structural organization of both inner and outer information processing circuits (see Clark 2016). Such forms of control may be important in terms of an agent's ability to deal with the training demands of a complex task (see Clowes and Morse 2005).

[15] Such systems are of particular interest given the aforementioned role of maturational processes in helping to shape the course of stage-dependent learning in the language domain (see Elman 1993).

number of different views as to what is implied 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 (Anderson 2003; Shapiro 2007, 2011). Research into embodied cognition thus typically focuses on 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 (e.g., the position of their eyes and ears) helps to structure incoming sensory information in ways that support perceptual processing (Webb 1996). Other forms of embodied cognition research seek to evaluate the role of physical actions (e.g., hand gestures) in supporting various forms of human cognitive competence (Goldin-Meadow 2003; Goldin-Meadow and Beilock 2010; Cook et al 2008; Andrade 2010)[16].

In the case of naturally-intelligent (i.e., biological) systems, the body would seem to be a natural focus of attention for those embracing a situated and ecologically-oriented approach to cognition. But do any of the insights gleaned from embodied cognitive science have any relevance to contemporary research on machine intelligence, especially when it comes to the species of intelligent systems that inhabit the ecological niche of the Internet? Such a claim is likely to be immediately rejected by those who insist that our intuitions as to what constitutes the 'body' is determined by issues of material composition. It might be said, therefore, that issues of embodied cognition are irrelevant to machine intelligence on the grounds that the body is a biological entity and only biological systems can be said to qualify as embodied systems. 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 role of biology are liable to shift once we reflect on the corporeal status of a variety of non-biological objects; for example, tooth implants, prosthetic limbs and cochlear implants. It should be clear, in this case, that the appeal to biology is probably insufficient in terms of determining what it is that makes something a part of the body. Instead, what seems to be important is the way in which some object or set of objects mediates our sensorimotor engagements

[16] There are actually two ways one can view the notion of embodied cognition. These views parallel the distinction between extended and embedded approaches to cognition. One can thus adopt an 'extended' view of embodied cognition in which (extra-neural) bodily elements are seen as part of the physical machinery that realizes mental states and processes. Alternatively, one can adopt a more 'embedded' view of embodied cognition in which (extra-neural) bodily elements are seen as influencing the shape of cognitive processes without thereby forming part of the causally-active physical web of material elements that realizes those processes. In practice, I suggest that is rather hard to avoid the extended view of embodied cognition. Firstly, there is a sense in which the body is an obvious constituent element in any (overt) intelligent performance. Secondly, once we reflect on the physiological role of glial cells in neural functioning (Araque and Navarrete 2010; Perea et al 2014; Oliveira et al 2015), it seems as though even the most ardent advocate of neurocentrism would be inclined to acknowledge the role of some *non-neural* bodily elements (e.g., astrocytes) in the mechanistic realization of cognitive phenomena, even those that, in this case, occur solely within the intra-cranial domain.

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 in which they service our locomotory objectives; and our teeth count as part of our body because of the way in which 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 robotics. It thus seems that some of the core lessons and ideas from embodied cognitive science are just as applicable to synthetic intelligent systems as they are to more natural (i.e., biological) ones. Consider, for example, some of the early work in robotics where adaptive behavioral success was seen to result from the exploitation of simple body-related sensorimotor contingencies. Clark's (1999) description of the can-collecting robot, Herbert, for example, highlights the role of bodily motion (in this case, the rotation of the robot's 'torso') in securing an appropriate physical alignment between the robot and a target object—one that helps to secure the success of subsequent reaching movements.

More recent work concerning the design of intelligent systems also highlights the value of embodied perspectives that are appropriately divested of bio-centric biases. Good examples, here, can be found in the work on morphological computation (Paul 2006; Pfeifer and Gómez 2009), where, as the name suggests, we encounter a direct appeal to the role of bodily structure in mediating or realizing various forms of computational processing. Work at the interface of robotics and developmental science has also revealed the potential value of an embodied perspective. Consider, for example, the aforementioned work by Gómez et al. (2004) showing how developmentally-informed alterations in sensory and motor capabilities can play a productive role in yielding enhanced forms of perceptuo-motor success in a real-world robot (see Section 5). In this case, a non-biological conception of the body seems *essential* to our understanding of what it is that is yielding elevated levels of performance.

I take it that most readers are, by now, likely to be convinced of the case for a non-biological conception of the 'body'. The question is: what is it that determines when some set of physical elements should qualify as a genuine body (or body part)? The answer to this question, I suggest, should be grounded in appeals to the functional contribution of the physical elements in (e.g.) mediating our sensorimotor engagements with the wider extra-corporeal environment. Something along these lines is suggested by Clark (2008). He 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, and the stable (though not permanently fixed) platform whose features and relations can be relied upon (without being represented) in the computations underlying some intelligent performances" (p. 207). The claim, in essence, is that we should identify the body with whatever it is that is playing the sort of role that our biological body typically plays with respect to

the genesis and organization of intelligent behavior. Once we adopt a functional approach to the conception of the 'body', however, then there is no reason why we should feel obliged to reject the idea of embodied forms of machine intelligence.

The question remains, however, as to whether the notion of embodied cognition has any substantive value in helping us to understand (and engineer) Internet-situated forms of machine intelligence. What, exactly, we might ask, counts as the body (however conceived) in the case of online intelligent systems? One answer to that question comes from recent work on what is referred to as the 'global brain' (Heylighen 2013; Kyriazis 2015; Vidal 2015). The global brain, in this case, is a metaphor for a form of planetary-scale superintelligence that is "formed by all people on this planet together with their artifacts and technologies" (Heylighen 2013, p. 905). The idea, in essence, is that global information and communication networks, such as the contemporary Internet, serve as (part of) the realization base for "problem-solving abilities...[that are] orders of magnitudes larger than that of any single individual, organization, or computer system" (Heylighen 2013, p. 905). Crucially (although somewhat confusingly given their appeal to neurological metaphors), global brain theorists accept the idea that intelligence is constituted by more than just the operation of 'brain-based' processes. They thus accept the claim, which lies at the heart of the theoretical agenda of embodied cognitive science, that corporeal elements can (and typically will) play functionally-significant roles in supporting the emergence of intelligent performances (see Heylighen 2012). This much is apparent in the way in which proponents of the global brain talk about the contribution of human and technological elements to 'embodied intelligence':

> "The accelerating incorporation of a multitude of hardware sensors and effectors (e.g. satellites, cameras, remote controls) together with an ever more intimate individual-net interface turn the global brain into a truly situated and embodied intelligence...Every human being or piece of machinery hooked up to the web forms part of its body, providing it with additional capabilities for input, output and control" (Heylighen 2012).

Here, then, we encounter one way in which issues of embodied cognition might be able to establish a useful foothold in discussions regarding the nature of Internet-based machine intelligence. The general idea is that, in at least some cases, it may be appropriate to see technological devices (and perhaps even human agents) as literal body parts, i.e., as physical elements that play the same sort of functional role—e.g., the mediation of real-world sensorimotor engagements—that we typically associate with *bona fide* body parts. Inasmuch as such claims are true, we are provided with an important opportunity to consider how issues of embodiment might affect the attempt to engineer online forms of machine intelligence. An important issue for future research is thus to determine whether some of the insights gleaned from research into biological (and robotic) systems can be applied to the realm of online systems and services. This, I suggest, should attempt to go beyond the rather simple

case of informational 'pickup' where some kind of object is used to provide a form of 'perceptual' contact with the real-world environment[17]. In fact, I suggest that the most compelling case for an embodied cognitive science of Internet-situated machine intelligence is likely to come about as a result of the attempt to allow machine-based systems to exert some form of control over their 'perceptual' environments. One example might be the case of sensor devices being actively moved and configured in order to structure the flow of incoming sensory information. Another might be the case of human subjects being actively recruited—perhaps via crowdsourcing techniques—to engage in tasks that involve the transformation of problematic bodies of target data (e.g., photorealistic images) into something more suitably aligned with the idiosyncratic processing capabilities of a symbol-crunching computational machine[18]. These sort of closed-loop sensorimotor engagements with some 'external' body of information strike at the heart of an important theme of embodied cognitive science, namely the way in which physical actions can sometimes be used to transform the nature of the computational task confronting the cognitive agent (Kirsh and Maglio 1994) (see also Section 4).

7 Collective Cognition

For many supporters of an ecological approach to cognition, the social environment emerges as a particularly important focus of theoretical and scientific attention (e.g., Hutchins 2010). This is evident in work relating to what is called distributed (Hutchins 1995a), collective (Smart et al in press), team (Cooke et al 2004), group (Theiner 2014; Theiner and O'Connor 2010), swarm (Trianni et al 2011; Turner 2011), and colony-level (Marshall and Franks 2009) cognition. The core idea, in this case, is that multiple individuals are involved (as constitutive elements) in the mechanistic realization of cognitive process-

[17] Some examples, here, include cases of participatory sensing (Carrera et al 2013; Resch 2013) and collective sensing (Dugas et al 2013; Ginsberg et al 2009; Resch 2013). In the case of collective sensing, for example, the inputs provided by human agents can sometimes be used to support the detection of phenomena that are otherwise very difficult to discern. A notable example, here, concerns the use of query-based syndromic surveillance techniques to monitor trends in disease pandemics. By monitoring the prevalence of semantically-relevant Google search terms, for example, it has proved possible to detect and monitor the progress of influenza outbreaks (Dugas et al 2013; Ginsberg et al 2009). Similarly, search query behaviour has yielded insights into the seasonal nature of mental health concerns (Ayers et al 2013), as well as information about the side-effects of particular medicines (White et al 2013).

[18] Such human-mediated transformations of the target data could assume a variety of forms. One example is provided by systems that support the (human-assisted) semantic annotation of complex resources, such as medical X-ray images (see Rubin et al 2008). These sorts of interventions are typically aimed at improving machine-based access to important bodies of task-relevant information that are otherwise 'hidden' or 'deeply embedded' in the focal resource. Human annotations, in this case, help to resolve what is sometimes referred to as the 'semantic gap' when it comes to the analysis of digital content (see Siorpaes and Simperl 2010).

ing routines that are distributed across the multi-agent ensemble. In essence, a community of individual agents is itself seen as a form of cognitive system[19].

There are clearly many situations in which the Internet can be seen to support human-based forms of collective cognition. We see this, for example, in research relating to social computation (Kearns 2012) and augmented social cognition (Chi 2009), as well as in work relating to issues of collective intelligence (Malone and Bernstein 2015), the global brain (Heylighen 2013) and human computation (Michelucci and Dickinson 2016). From the perspective of machine intelligence, however, the Internet can also be seen to provide an important form of support for collective cognizing. By this I do not mean that the Internet simply provides the basis for the emergence of novel kinds of socio-technical systems—online versions of the socio-technical systems that are the typical focus of attention by distributed cognition theorists (Hutchins 1995a,b). This much is, of course, trivially true, in the sense that every instance of Internet-enabled collective cognition will invariably involve *some* form of technological mediation, even when the focus of attention is on a collection of human agents. Instead, what I mean by the idea of the Internet providing support for machine-based forms of collective cognition is that the Internet, *qua* global information and communication network, provides the basis for cognitively-potent forms of 'social' interaction and engagement between synthetic agents. The Internet is thus a potential boon for researchers who work in the areas of distributed computation and distributed artificial intelligence (see Clearwater et al 1992): it enables an episode of machine-based processing to be distributed across multiple independent (synthetic) agents.

Of course, the ability to participate in episodes of collective cognizing requires an ability to communicate[20], and it is here, I suspect, that we are likely to see important advances in machine-based capabilities. One area of research that is likely to prove of particular interest relates to the attempt to equip machines with language generation capabilities (Wright 2015). At the present time, such capabilities are intended to support the generation of text that is intended for human consumption. In the future, however, it is possible that the primary beneficiaries of these machine-based narrative ca-

[19] Such system-oriented views often appeal to the idea that *all* cognition is in some sense distributed (see Hutchins, 2014) and evident at multiple organizational levels. Turner (2011), for example, notes that: "Brains and sensory systems are the usual context for our thinking about cognitive systems, but there is no inherent reason why cognitive systems need be circumscribed in this way. Life occurs at multiple scales of organization, ranging from cell to organism to superorganism. If cognition is useful for one level of organization, as it clearly is for organisms, it is difficult to argue that cognition would not be useful, or possible, at higher organizational levels as well" (p. 42).

[20] In the absence of communicative abilities, no amount of physical networking (e.g., electronic cabling) will suffice to support collective cognition. It is in this sense that we can perhaps see the development of communicative capabilities as akin to a form of cognitive networking capability—one that depends on, but is nevertheless distinct from, the underlying physical network. The same is, of course, true of human agents. As noted by Donald (1991), "Individuals in possession of reading, writing, and other visuographic skills thus become somewhat like computers with networking capabilities; they are equipped to interface, to plug into whatever network becomes available" (p. 311).

pabilities will be other language-enabled machine agents. This is a possibility that is raised by Peter Swirski (2013). He suggests that eventually the works of computational authors (or 'computhors') might be 'read' solely by other computational agents. The result is that we would see the emergence of a new kind of linguistic community, one that is entirely devoid of human involvement: machines generating linguaform content solely for the cognitive benefit of other machines.

It is at this point that a number of important issues come into clearer focus. One issue concerns the extent to which we may encounter a form of evolutionary process that alters the nature of the linguistic vehicles that are used to communicate information. In this case, it is important to take note of a body of work that sees linguistic vehicles as subject to iterative processes of trans-generational modification. It is via these transformational process that communicative resources and conventions may become adaptively aligned with pre-existing agential abilities (Christiansen and Kirby 2003; Chater and Christiansen 2010), enhancing the extent to which the agent community is able to achieve collective forms of problem-solving success (see Hutchins and Hazelhurst 1991). Another idea relates to what is sometimes referred to as the supracommunicative view of language: the idea that language helps to transform the cognitive capabilities of the language-wielding agent (Clark 1998, 2012). Finally, we should not lose sight of the fact that in generating informational content—much of which may only be comprehensible to synthetic agents—machine-based systems can play an important role in structuring the environment that informs the shape of their existing cognitive and epistemic capabilities. This is an idea that, as we shall see in the next section, dovetails perfectly with ideas concerning ecological engineering (Sterelny 2003) and cognitive niche construction (Clark 2008). For present purposes, the idea of machines actively structuring and configuring aspects of the online environment should strike a chord with those who are familiar with the notions of swarm intelligence, swarm cognition and stigmergic algorithms[21] (Garnier et al 2007; Trianni et al 2011). As is apparent from the study of comparatively simple animals (e.g., insects) the ability to actively structure the environment in particular ways (e.g., via the use of pheromone trails) can often be used to achieve various forms of collective cognitive success (Couzin 2009), and in many cases, these kinds of collective accomplishment can seem profoundly out of kilter with the cognitive abilities of the individual agents who comprise the collective organization (e.g., Franks et al 2002). In future research, it will be important to determine the extent to which similar forms of (surprising) collective cognitive success can emerge from communities of machine agents that are equipped with the means to engage in both direct and indirect (i.e., stigmergic) forms of inter-agent communication.

[21] Stigmergy is a form of social influence that occurs as the result of the modification of the local environment. A useful definition is provided by Heylighen (2016): "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).

8 Ecological Engineering

Throughout this paper, I have sought to draw on the idea that the Internet forms an important part of the cognitive ecology for future forms of machine intelligence. There are a variety of reasons to think that this sort of ecological approach to the Internet is useful, not least because it helps to support the kinds of analyses that were undertaken in respect of extended, embedded, scaffolded, embodied and collective cognition. Now that these analyses have been undertaken, it is time to highlight an unexpected bonus of the ecological perspective. This bonus comes into view once we pay attention to the emphasis that is placed on the notions of ecological engineering (Sterelny 2003) and cognitive niche construction (Clark 2008). The core idea, in brief, is that biological agents can sometimes play an active role in creating and configuring the environments that subsequently work to shape, scaffold and support their cognitive processing routines. In defining the notion of cognitive niche construction, for example, Clark (2008) suggests that we should view it 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). Such processes are deemed to be important because they provide a powerful means of bootstrapping certain kinds of cognitive capability. In the case of extended cognition, for example, we humans have the opportunity to create the very resources that are then factored deep into our cognitive processing routines. We thus have to the ability to engage in an important form of cognitive 'self-engineering', effectively creating the conditions under which we are able to the extend the reach of our cognitive capabilities. "We use intelligence," writes Clark, "to structure our environment so that we can succeed with less intelligence. Our brains make the world smart so that we can be dumb in peace!" (Clark 1997, p. 180).

Perhaps unsurprisingly, the notion of cognitive niche construction is just as applicable to the realm of the online environment as it is to the realm of the offline environment. This is particularly so, once we consider the nature of our current interaction with the Internet. With the advent of what is sometimes referred to as Web 2.0, we have witnessed a transition in the way human agents interact with the Web. Instead of merely consuming online content, human agents are now able to create and configure the resources that make up the environment of the online world. This is perhaps most evident in the case of systems, such as Wikipedia, where the vast majority of content is generated by the human user community. Other prominent examples include social networking sites, such as Facebook, microblogging services, such as Twitter, and social media sites, such as YouTube and Flickr[22]. Inasmuch as these systems can be seen as part of the cognitive ecology of the Internet, then it seems appropriate to see ourselves (both individually and collectively) as engaged in a

[22] These systems are sometimes referred to as 'social machines' (see Hendler and Berners-Lee 2010). The role of the human user community in generating or creating online content is sometimes seen as the defining feature of such systems (Hendler and Berners-Lee 2010, p. 156).

process of online ecological engineering or online cognitive niche construction, actively creating and configuring the environment that influences the profile of our subsequent cognitive and epistemic endeavours.

But it is not only human cognition that stands to be affected by this sort of ecological engineering. If the arguments rehearsed in this paper are anywhere near the mark, then it becomes clear that we should see the actions of human agents as not only altering the ecological setting for human intelligence, but as also affecting the environment in which various forms of machine intelligence are materially embedded. We thus begin to creep up on an idea that first surfaced in the discussion of scaffolded cognition (see Section 5). This is the idea that the Internet affords a relatively unique opportunity for the global human community to create the sorts of conditions under which advanced forms of machine intelligence might emerge. Contrary to any radical sea-change in our approach to machine learning, or some dramatic breakthrough in our understanding of artificial intelligence, perhaps what is really required to advance the state-of-the-art in machine intelligence is merely(!) a means of creating the sort of ecological niche that enables machine-based cognitive systems to thrive and flourish. Our online interactions, if this is correct, might be seen as a means of progressively shaping (and re-shaping) the specific mix of opportunities and affordances that help to bring the next generation of cognitive computing systems into being.

If all this still seems somewhat vague, perhaps it will help to consider a concrete example of how human activity, when appropriately supported and enabled by the technological infrastructure of the Internet, can lead to advances in machine intelligence. Take IBM Watson, the poster child of the emerging cognitive computing paradigm (Kelly and Hamm 2013). Watson's virtuoso performance in answering of array of difficult questions is undoubtedly an important demonstration of the growing sophistication of machine-based processing, especially in the areas of natural language processing, machine learning and domain-specific reasoning. However, in the course of being awestruck by (at least some of) Watson's outputs, it is easy to overlook the simple fact that many of the inputs to the system—especially the information resources that Watson exploits in supplying its human interlocutors with answers—are ones that are, in general, developed by large numbers of human individuals. Such resources include online encyclopedias, dictionaries, thesauri, taxonomies, ontologies and so on (Ferrucci et al 2010). None of these resources, it should be clear, were specifically intended to support Watson, or indeed any of the other cognitive computing systems that are the focus of current research attention. Rather, what has happened is that, as a result of the Internet's ability to support and solicit online engagement, we now have access to an array of resources that can be used to fuel the development of relatively novel forms of machine intelligence. If such systems—I am thinking primarily of Watson here—had appeared in the absence of the Internet, we would, I suspect, have been astonished at their capabilities. And yet Watson is, at root, exactly the same kind of symbol manipulating device that has long been been the focus of philosophical theorizing and the primary instrument of cognitive scientific

practice. Arguably what has changed in recent times is not so much the underlying technology, as the nature of the environment in which such technologies are situated. The Internet has thus enabled us to create the sort of environment in which a conventional symbol crunching computational economy is able to exhibit forms of intelligence that might have seemed utterly out of reach just a few short years ago.

The key take home message, here, is that it sometimes makes sense to see intelligence as (in part) a product of the environment in which an agent is materially embedded. Inasmuch as this is true, and inasmuch as it applies to the case of machine intelligence, then it clearly makes sense to see issues of environmental structuring, enrichment and configuration as of crucial importance to the enhancement of machine-based cognitive capabilities. Given that the human-engineered environment is deemed to be of crucial relevance to our own human cognitive capabilities (Clark 2003, 2008), why not assume that the key to advanced forms of machine intelligence is likewise to be found in the ecology wherein a growing number of our intelligent machines are now situated?

9 Conclusion

In the present paper, I have adopted the idea that the Internet is a form of cognitive ecology. The primary value of this perspective is that it helps us gain a better understanding of how the Internet might shape and influence the profile of machine-based cognitive capabilities, both now and in the future. The idea that the Internet forms part of the cognitive ecology for machine intelligence also brings a range of other benefits. For example, it enables us to establish a useful point of contact with work that emphasizes an ecological approach to *human* cognition (Hutchins 2010; Malafouris 2013; Tribble and Sutton 2011; Bateson 1972). It should also be clear that an ecological approach enables us to take a number of different conceptual stances when it comes to an analysis of the cognitive significance of the Internet. We can thus view the Internet from the perspective of those approaches that are typically associated with the situated cognition movement, e.g., extended, embodied and collective cognition (see Robbins and Aydede 2009b). As we have seen, when we apply this situated perspective to the case of Internet-based forms of machine intelligence, a number of interesting and important issues come to light. These include, but are not necessarily limited to, the following:

– **Philosophical Issues:** As we saw in the discussion relating to HEMC, a situated approach to machine intelligence can help to reshape and reorient the trajectory of existing philosophical debates. When it comes to extended cognition, for example, much of the philosophical discourse is based around the idea of human-centered—or organism-centered (see Clark 2008)—cognition. By applying the notion of extended cognition to the realm of machine intelligence, we help to broaden the philosophical debate and (hopefully) instigate new forms of philosophical progress.

- **Design and Development Issues:** A situated approach is also useful in helping to inform the design and development of online systems. We saw, for example, that manually annotated resources can sometimes be used to support the development of fully automated capabilities (see Section 5). Once we appreciate this possibility, it alters the way we approach the design and development of subsequent systems. In the context of the discussion on collaborative filtering (see Section 4), we also encountered the idea that the human social environment could be exploited as a form of computationally-significant resource. This is an idea that seems to be gaining increasing traction within the computer science community (Bernstein et al 2012; Robertson and Giunchiglia 2013).
- **Research Issues:** Finally, a situated approach to the Internet reveals a number of new directions for scientific research. These include, for example, the idea that maturationally-relevant changes in sensory, motor and cognitive resources, coupled with the sort of scaffolding provided by the human social environment, might help to advance the current state-of-the-art in machine intelligence (see Section 6).

Perhaps what is most important about the idea of the Internet as a form of cognitive ecology is that it helps us to appreciate the transformative potential of the Internet vis-à-vis future forms of machine intelligence. Given the way in which an ecological perspective encourages us to see cognition as heavily dependent on, if not partly constituted by, aspects of the wider environment, we can view the Internet as marking a potentially significant waypoint in the 'evolution' of machine-based capabilities. In order to help us see this we need only cast our eye towards digital global positioning systems. The constant stream of data provided by the Global Positioning System (GPS) has obviously influenced the way in which we humans navigate and locate ourselves in space. It should also be clear, however, that the *very same system* has had a significant impact on the development of spatially-aware autonomous systems. The availability of GPS signals has thus transformed the kinds of approach that can be adopted with respect to the implementation of (e.g.) driverless cars and pilotless drones, enabling forms of navigational competence that may have been difficult, if not impossible, to engineer in the absence of a suitably structured and enriched ambient environment.

What is true of one kind of digital data ecology (i.e., GPS signals) is also, I suggest, true of another kind of digital data ecology (i.e., the Internet). But the Internet has properties that seem to be of particular relevance and importance to machine intelligence. We have seen, for example, how the Internet provides an unprecedented opportunity to tap into human capabilities in a manner and on a scale that would have been unthinkable just a couple of decades ago (see Section 3). We have also encountered the idea that the Internet provides access to a variety of kinds of training data that are progressively altering our traditional approaches to machine learning (see Section 5). Finally, we have seen how the Internet affords access to bodies of information that can be used to develop intelligent systems whose capabilities are, at least in some respects,

superior to those of their human counterparts (recall the discussion of IBM Watson in Section 8).

The advent of the Internet thus marks a potentially important milestone in our attempts to engineer a range of relatively novel forms of intelligent system. In helping us to press maximal cognitive benefit from this new cognitive ecology, we should, I suggest, seek to establish intellectual contact with work in the philosophy of mind that aims to illuminate the ways in which human cognition is influenced by an array of extra-neural and extra-organismic resources. Such points of contact may be invaluable in terms of helping us to appreciate the many ways in which the Internet helps to shape the intelligence of both ourselves and the machines we build.

Acknowledgements 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.

References

Adams F, Aizawa K (2010) Defending the bounds of cognition. In: Menary R (ed) The Extended Mind, MIT Press, Cambridge, Massachusetts, USA
von Ahn L (2006) Games with a purpose. Computer 39(6):96–98
von Ahn L, Dabbish L (2008) Designing games with a purpose. Communications of the ACM 51(8):58–67
Amatriain X (2012) Mining large streams of user data for personalized recommendations. ACM SIGKDD Explorations Newsletter 14(2):37–48
Anderson M (2003) Embodied cognition: A field guide. Artificial Intelligence 149(1):91–130
Andrade J (2010) What does doodling do? Applied Cognitive Psychology 24(1):100–106
Araque A, Navarrete M (2010) Glial cells in neuronal network function. Philosophical Transactions of the Royal Society B: Biological Sciences 365(1551):2375–2381
Ayers JW, Althouse BM, Allem JP, Rosenquist JN, Ford DE (2013) Seasonality in seeking mental health information on Google. American Journal of Preventive Medicine 44(5):520–525
Barrett L (2011) Beyond the Brain: How Body and Environment Shape Animal and Human Minds. Princeton University Press, Princeton, New Jersey, USA
Barrington L, Turnbull D, Lanckriet G (2012) Game-powered machine learning. Proceedings of the National Academy of Sciences 109(17):6411–6416
Barsalou L (2008) Grounded cognition. Annual Review of Psychology 59:617–645
Barsalou LW (2010) Grounded cognition: Past, present, and future. Topics in Cognitive Science 2(4):716–724
Bateson G (1972) Steps to an Ecology of Mind. Jason Aronson Inc., London, UK
Berners-Lee T, Hendler J, Lassila O (2001) The Semantic Web. Scientific American 284(4):34–43
Bernstein A, Klein M, Malone TW (2012) Programming the global brain. Communications of the ACM 55(5):41–43
Bjorklund DF (1997) The role of immaturity in human development. Psychological Bulletin 122(2):153–169
Branson S, Horn G, Wah C, Perona P, Belongie S (2014) The ignorant led by the blind: A hybrid human-machine vision system for fine-grained categorization. International Journal of Computer Vision 108(1):3–29
Brin S, Page L (1998) The anatomy of a large-scale hypertextual search engine. In: 7th Annual Conference of the World Wide Web, Brisbane, Australia

Carrera F, Guerin S, Thorp J (2013) By the people, for the people: The crowdsourcing of "STREETBUMP": An automatic pothole mapping app. International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences XL-4/W1:19–23
Chater N, Christiansen MH (2010) Language evolution as cultural evolution: How language is shaped by the brain. Wiley Interdisciplinary Reviews: Cognitive Science 1(5):623–628
Chi EH (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
Christiansen MH, Kirby S (2003) Language evolution: Consensus and controversies. Trends In Cognitive Sciences 7(7):300–307
Clark A (1997) Being There: Putting Brain, Body and World Together Again. MIT Press, Cambridge, Massachusetts, USA
Clark A (1998) Magic words: How language augments human computation. In: Carruthers P, Boucher J (eds) Language and Thought: Interdisciplinary Themes, Cambridge University Press, Cambridge, UK
Clark A (1999) Where brain, body, and world collide. Cognitive Systems Research 1(1):5–17
Clark A (2003) Natural-Born Cyborgs: Minds, Technologies and the Future of Human Intelligence. Oxford University Press, Oxford, UK
Clark A (2007) Curing cognitive hiccups: A defense of the extended mind. Journal of Philosophy 104(4):163–192
Clark A (2008) Supersizing the Mind: Embodiment, Action, and Cognitive Extension. Oxford University Press, New York, New York, USA
Clark A (2010a) Coupling, constitution and the cognitive kind: A reply to Adams and Aizawa. In: Menary R (ed) The Extended Mind, MIT Press, Cambridge, Massachusetts, USA
Clark A (2010b) Much ado about cognition. Mind 119(476):1047–1066
Clark A (2012) How to qualify for a cognitive upgrade: Executive control, glass ceilings and the limits of simian success. In: McFarland D, Stenning K, McGonigle-Chalmers M (eds) The Complex Mind: An Interdisciplinary Approach, Palgrave Macmillan, Basingstoke, England, UK
Clark A (2016) Surfing Uncertainty: Prediction, Action and the Embodied Mind. Oxford University Press, New York, New York, USA
Clark A, Chalmers D (1998) The extended mind. Analysis 58(1):7–19
Clearwater SH, Hogg T, Huberman BA (1992) Cooperative problem solving. In: Huberman BA (ed) Computation: The Micro and the Macro View, World Scientific Publishing, London, UK
Clowes R (2015) Thinking in the cloud: The cognitive incorporation of cloud-based technology. Philosophy & Technology 28(2):261–296
Clowes RW, Morse AF (2005) Scaffolding cognition with words. In: Fifth International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems, Lund, Sweden
Coburn C (2014) Play to cure: Genes in space. Lancet Oncology 15(7):688
Cook SW, Mitchell Z, Goldin-Meadow S (2008) Gesturing makes learning last. Cognition 106(2):1047–1058
Cooke NJ, Gorman JC, Rowe LJ (2004) An ecological perspective on team cognition. In: Salas E, Goodwin GF, Burke CS (eds) Team Effectiveness in Complex Organizations: Cross-Disciplinary Perspectives and Approaches, Psychology Press, New York, New York, USA
Cooper S, Khatib F, Treuille A, Barbero J, Lee J, Beenen M, Leaver-Fay A, Baker D, Popović Z, Foldit Players (2010a) Predicting protein structures with a multiplayer online game. Nature 466(7307):756–760
Cooper S, Treuille A, Barbero J, Leaver-Fay A, Tuite K, Khatib F, Snyder AC, Beenen M, Salesin D, Baker D, Popović Z, Foldit Players (2010b) The challenge of designing scientific discovery games. In: 5th international Conference on the Foundations of Digital Games, Monterey, California, USA
Couzin I (2009) Collective cognition in animal groups. Trends In Cognitive Sciences 13(1):36–43

Dieleman S, Willett KW, Dambre J (2015) Rotation-invariant convolutional neural networks for galaxy morphology prediction. Monthly Notices of the Royal Astronomical Society 450(2):1441–1459

Donald M (1991) Origins of the Modern Mind: Three Stages in the Evolution of Culture and Cognition. Harvard University Press, Cambridge, Massachussets, USA

Dugas AF, Jalalpour M, Gel Y, Levin S, Torcaso F, Igusa T, Rothman RE (2013) Influenza forecasting with Google flu trends. PLoS One 8(2):e56176

Elman JL (1993) Learning and development in neural networks: The importance of starting small. Cognition 48(1):71–99

Engel AK, Friston KJ, Kragic D (eds) (2016) The Pragmatic Turn: Toward Action-Oriented Views in Cognitive Science. MIT Press, Cambridge, Massachusetts, USA

Ferrucci D, Brown E, Chu-Carroll J, Fan J, Gondek D, Kalyanpur AA, Lally A, Murdock JW, Nyberg E, Prager J (2010) Building Watson: An overview of the DeepQA project. AI Magazine 31(3):59–79

Fogarty J, Tan D, Kapoor A, Winder S (2008) CueFlik: Interactive concept learning in image search. In: SIGCHI Conference on Human Factors in Computing Systems, Florence, Italy

Franks NR, Pratt SC, Mallon EB, Britton NF, Sumpter DJ (2002) Information flow, opinion polling and collective intelligence in house-hunting social insects. Philosophical Transactions of the Royal Society of London B: Biological Sciences 357(1427):1567–1583

Froese T, Paolo D, Ezequiel A (2011) The enactive approach: Theoretical sketches from cell to society. Pragmatics & Cognition 19(1):1–36

Garnier S, Gautrais J, Theraulaz G (2007) The biological principles of swarm intelligence. Swarm Intelligence 1(1):3–31

Ginsberg J, Mohebbi MH, Patel RS, Brammer L, Smolinski MS, Brilliant L (2009) Detecting influenza epidemics using search engine query data. Nature 457(7232):1012–1014

Goldin-Meadow S (2003) Hearing Gesture: How Our Hands Help Us Think. Harvard University Press, Cambridge, Massachusetts, USA

Goldin-Meadow S, Beilock SL (2010) Action's influence on thought: The case of gesture. Perspectives on Psychological Science 5(6):664–674

Gómez G, Lungarella M, Eggenberger Hotz P, Matsushita K, Pfeifer R (2004) Simulating development in a real robot: On the concurrent increase of sensory, motor, and neural complexity. In: Fourth International Workshop on Epigenetic Robotics, Genoa, Italy

Good BM, Su AI (2011) Games with a scientific purpose. Genome Biology 12(135):1–3

Halpin H (2013) Does the Web extend the mind? In: 5th Annual ACM Web Science Conference, Paris, France

Halpin H, Clark A, Wheeler M (2010) Towards a philosophy of the Web: Representation, enaction, collective intelligence. In: Web Science Conference, Raleigh, North Carolina, USA

Hendler J, Berners-Lee T (2010) From the Semantic Web to social machines: A research challenge for AI on the World Wide Web. Artificial Intelligence 174:156–161

Heylighen F (2012) A brain in a vat cannot break out: Why the singularity must be extended, embedded and embodied. Journal of Consciousness Studies 19(1-2):126–142

Heylighen F (2013) From human computation to the global brain: The self-organization of distributed intelligence. In: Michelucci P (ed) Handbook of Human Computation, Springer, New York, New York, USA

Hirose N (2002) An ecological approach to embodiment and cognition. Cognitive Systems Research 3(3):289–299

Horvitz EJ (2007) Reflections on challenges and promises of mixed-initiative interaction. AI Magazine 28(2):19–22

Huang Y, Erdogmus D, Pavel M, Mathan S, Hild KE (2011) A framework for rapid visual image search using single-trial brain evoked responses. Neurocomputing 74(12):2041–2051

Hurwitz JS, Kaufman M, Bowles A (2015) Cognitive Computing and Big Data Analytics. John Wiley & Sons, Inc., Indianapolis, Indiana, USA

Hutchins E (1995a) Cognition in the Wild. MIT Press, Cambridge, Massachusetts, USA

Hutchins E (1995b) How a cockpit remembers its speeds. Cognitive Science 19(3):265–288

Hutchins E (2001) Distributed cognition. In: Smelser N, Baltes P (eds) International Encyclopedia of the Social & Behavioral Sciences, Elsevier Science Ltd, Oxford, UK

Hutchins E (2010) Cognitive ecology. Topics in Cognitive Science 2(4):705–715
Hutchins E, Hazelhurst B (1991) Learning in the cultural process. In: Langdon CG, Taylor C, Farmer JD, Rasmussen S (eds) Artificial Life II, Addison-Wesley Publishing Company, Redwood City, California, USA
Jacucci G, Gamberini L, Freeman J, Spagnolli A (eds) (2014) Symbiotic Interaction. Springer, London, UK
Kapoor A, Tan D, Shenoy P, Horvitz E (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
Kearns M (2012) Experiments in social computation. Communications of the ACM 55(10):56–67
Kelly JE, Hamm S (2013) Smart Machines: IBM's Watson and the Era of Cognitive Computing. Columbia University Press, Chichester, West Sussex, UK
Kendon V, Sebald A, Stepney S (2015) Heterotic computing: Past, present and future. Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences 373:20140225
Khatib F, Cooper S, Tyka MD, Xu K, Makedon I, Popović Z, Baker D, Foldit Players (2011a) Algorithm discovery by protein folding game players. Proceedings of the National Academy of Sciences 108(47):18,949–18,953
Khatib F, DiMaio F, Cooper S, Kazmierczyk M, Gilski M, Krzywda S, Zabranska H, Pichova I, Thompson J, Popović Z, Jaskolski M, Baker D (2011b) Crystal structure of a monomeric retroviral protease solved by protein folding game players. Nature Structural & Molecular Biology 18(10):1175–1177
Kirsh D (1995) The intelligent use of space. Artificial Intelligence 73(1-2):31–68
Kirsh D, Maglio P (1994) On distinguishing epistemic from pragmatic action. Cognitive Science 18:513–549
Kraut R, Maher ML, Olson J, Malone TW, Pirolli P, Thomas JC (2010) Scientific foundations: A case for technology-mediated social-participation theory. Computer 43(11):22–28
Kyriazis M (2015) Systems neuroscience in focus: From the human brain to the global brain? Frontiers in Systems Neuroscience 9(7):1–4
Lasecki WS, Bigham JP (2013) Interactive crowds: Real-time crowdsourcing and crowd agents. In: Michelucci P (ed) Handbook of Human Computation, Springer, New York, New York, USA
Law E, von Ahn L (2011) Human computation. Synthesis Lectures on Artificial Intelligence and Machine Learning 5(3):1–121
Licklider JCR (1960) Man-computer symbiosis. IRE Transactions on Human Factors in Electronics HFE-1(1):4–11
Linden G, Smith B, York J (2003) Amazon.com recommendations: Item-to-item collaborative filtering. IEEE Internet Computing 7(1):76–80
Lintott CJ, Reed J (2013) Human computation in citizen science. In: Michelucci P (ed) Handbook of Human Computation, Springer, New York, New York, USA
Lintott CJ, Schawinski K, Slosar A, Land K, Bamford S, Thomas D, Raddick MJ, Nichol RC, Szalay A, Andreescu D, Murray P, van den Berg J (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
Malafouris L (2013) How Things Shape the Mind: A Theory of Material Engagement. MIT Press, Cambridge, Massachusetts, USA
Malone TW, Bernstein MS (eds) (2015) Handbook of Collective Intelligence. MIT Press, Cambridge, Massachusetts, USA
Marshall JAR, Franks NR (2009) Colony-level cognition. Current Biology 19(10):R395–R396
Marx V (2013) Neuroscience waves to the crowd. Nature Methods 10(11):1069–1074
Michelucci P (ed) (2013) Handbook of Human Computation. Springer, New York, New York, USA
Michelucci P (2016) Human computation and convergence. In: Bainbridge W, Roco M (eds) Handbook of Science and Technology Convergence, Springer, New York, New York, USA
Michelucci P, Dickinson JL (2016) The power of crowds. Science 351(6268):32–33

Modha DS, Ananthanarayanan R, Esser SK, Ndirango A, Sherbondy AJ, Singh R (2011) Cognitive computing. Communications of the ACM 54(8):62–71

Nagai Y, Kawai Y, Asada M (2011) Emergence of mirror neuron system: Immature vision leads to self-other correspondence. In: IEEE International Conference on Development and Learning, Frankfurt, Germany

Najafabadi MM, Villanustre F, Khoshgoftaar TM, Seliya N, Wald R, Muharemagic E (2015) Deep learning applications and challenges in big data analytics. Journal of Big Data 2(1):1–21

Neisser U (1997) The ecological study of memory. Philosophical Transactions of the Royal Society B: Biological Sciences 352(1362):1697–1701

Newport EL (1990) Maturational constraints on language learning. Cognitive Science 14(1):11–28

O'Leary DE (2013) Artificial intelligence and big data. IEEE Intelligent Systems 28(2):96–99

Oliveira JF, Sardinha VM, Guerra-Gomes S, Araque A, Sousa N (2015) Do stars govern our actions? Astrocyte involvement in rodent behavior. Trends in Neurosciences 38(9):535–549

Paul C (2006) Morphological computation: A basis for the analysis of morphology and control requirements. Robotics and Autonomous Systems 54:619–630

Pecher D, Zwaan R (eds) (2005) Grounding Cognition: The Role of Perception and Action in Memory, Language and Memory. Cambridge University Press, Cambridge, UK

Perea G, Sur M, Araque A (2014) Neuron-glia networks: Integral gear of brain function. Frontiers in Cellular Neuroscience 8(378):1–9

Pezzulo G, Barsalou LW, Cangelosi A, Fischer MH, McRae K, Spivey MJ (2013) Computational grounded cognition: A new alliance between grounded cognition and computational modeling. Frontiers in Psychology 3(612):1–11

Pfeifer R, Gómez G (2009) Morphological computation—Connecting brain, body, and environment. In: Sendhoff B, Körner E, Sporns O, Ritter H, Doya K (eds) Creating Brain-Like Intelligence: From Basic Principles to Complex Intelligent Systems, Springer, Berlin, Germany

Resch B (2013) People as sensors and collective sensing-contextual observations complementing geo-sensor network measurements. In: Krisp JM (ed) Progress in Location-Based Services, Springer, Berlin, Germany

Robbins P, Aydede M (eds) (2009a) The Cambridge Handbook of Situated Cognition. Cambridge University Press, New York, New York, USA

Robbins P, Aydede M (2009b) A short primer on situated cognition. In: Robbins P, Aydede M (eds) Cambridge Handbook of Situated Cognition, Cambridge University Press, Cambridge, UK

Robertson D, Giunchiglia F (2013) Programming the social computer. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences 371(1987):20120379

Rowlands M (2006) Body Language: Representation in Action. MIT Press, Cambridge, Massachusetts, USA

Rubin DL, Rodriguez C, Shah P, Beaulieu CF (2008) iPad: Semantic annotation and markup of radiological images. In: AMIA 2008 Annual Symposium, Washington DC, USA

Rupert R (2004) Challenges to the hypothesis of extended cognition. Journal of Philosophy 101(8):389–428

Rupert R (2009) Cognitive Systems and the Extended Mind. Oxford University Press, New York, New York, USA

Savage N (2012) Gaining wisdom from crowds. Communications of the ACM 55(3):13–15

Settles B (2012) Active learning. Synthesis Lectures on Artificial Intelligence and Machine Learning 6(1):1–114

Shapiro L (2007) The embodied cognition research programme. Philosophy Compass 2(2):338–346

Shapiro LA (2011) Embodied Cognition. Routledge, Abingdon, Oxon, UK

Shapiro LA (ed) (2014) The Routledge Handbook of Embodied Cognition. Routledge, New York, New York, USA

Siorpaes K, Simperl E (2010) Human intelligence in the process of semantic content creation. World Wide Web 13(1-2):33–59

Smart PR (2012) The Web-extended mind. Metaphilosophy 43(4):426–445
Smart PR (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
Smart PR (2014) Embodiment, Cognition and the World Wide Web. In: Shapiro LA (ed) The Routledge Handbook of Embodied Cognition, Routledge, New York, New York, USA
Smart PR (in press) Emerging digital technologies: Implications for extended conceptions of cognition and knowledge. In: Carter AJ, Clark A, Kallestrup J, Palermos OS, Pritchard D (eds) Extended Epistemology, Oxford University Press, Oxford, UK
Smart PR, Shadbolt NR (2014) Social machines. In: Khosrow-Pour M (ed) Encyclopedia of Information Science and Technology, IGI Global, Hershey, Pennsylvania, USA
Smart PR, Huynh TD, Braines D, Shadbolt NR (2010) Dynamic networks and distributed problem-solving. In: Knowledge Systems for Coalition Operations (KSCO'10), Vancouver, British Columbia, Canada
Smart PR, Heersmink R, Clowes R (in press) The cognitive ecology of the Internet. In: Cowley SJ, Vallée-Tourangeau F (eds) Cognition Beyond the Brain, 2nd edn, Springer-Verlag, London, England, UK
Staley DJ (2014) Brain, Mind and Internet: A Deep History and Future. Palgrave Macmillan, Basingstoke, England, UK
Sterelny K (2003) Thought in a Hostile World: The Evolution of Human Cognition. Blackwell Publishing, Oxford, UK
Sterelny K (2010) Minds: Extended or scaffolded? Phenomenology and the Cognitive Sciences 9(4):465–481
Stewart JR, Gapenne O, Di Paolo EA (2010) Enaction: Toward a New Paradigm for Cognitive Science. MIT Press, Cambridge, Massachusetts, USA
Sutton J (2010) Exograms and interdisciplinarity: History, the extended mind, and the civilizing process. In: Menary R (ed) The Extended Mind, MIT Press, Cambridge, Massachusetts, USA
Swirski P (2013) From Literature to Biterature: Lem, Turing, Darwin, and Explorations in Computer Literature, Philosophy of Mind, and Cultural Evolution. McGill-Queens University Press, Montreal, Canada
Taraborelli D (2008) How the Web is changing the way we trust. In: Waelbers K, Briggle A, Brey P (eds) Current Issues in Computing and Philosophy, IOS Press, Amsterdam, The Netherlands
Theiner G (2014) Varieties of group cognition. In: Shapiro LA (ed) The Routledge Handbook of Embodied Cognition, Routledge, New York, New York, USA
Theiner G, O'Connor T (2010) The emergence of group cognition. In: Corradini A, O'Connor T (eds) Emergence in Science and Philosophy, Routledge, Abingdon, Oxon, UK
Theiner G, Allen C, Goldstone RL (2010) Recognizing group cognition. Cognitive Systems Research 11:378–395
Trianni V, Tuci E, Passino KM, Marshall JA (2011) Swarm cognition: An interdisciplinary approach to the study of self-organising biological collectives. Swarm Intelligence 5(1):3–18
Tribble E, Sutton J (2011) Cognitive ecology as a framework for Shakespearean studies. Shakespeare Studies 39:94–103
Turner JS (2011) Termites as models of swarm cognition. Swarm Intelligence 5(1):19–43
Vidal C (2015) Distributing cognition: From local brains to global brain. In: Goertzel B, Goertzel T (eds) The End of the Beginning: Life, Society and Economy on the Brink of the Singularity, Humanity+ Press, Los Angeles, California, USA
Vygotsky LS (1986) Thought and Language. MIT Press, Cambridge, Massachusetts, USA
Webb B (1996) A cricket robot. Scientific American 275(6):94–99
White RW, Tatonetti NP, Shah NH, Altman RB, Horvitz E (2013) Web-scale pharmacovigilance: Listening to signals from the crowd. Journal of the American Medical Informatics Association 20(3):404–408
Wilson RA, Clark A (2009) Situated cognition: Letting nature take its course. In: Robbins P, Aydede M (eds) Cambridge Handbook of Situated Cognition, Cambridge University Press, Cambridge, UK

Wright A (2015) Algorithmic authors. Communications of the ACM 58(11):12–14