The Extended Mind and Network-Enabled Cognition

Paul R. Smart[1], Paula C. Engelbrecht[2], Dave Braines[3], James A. Hendler[4] and Nigel Shadbolt[1]

[1]School of Electronics and Computer Science, University of Southampton, Southampton, SO17 1BJ, UK.
[2]School of Psychology, University of Southampton, Southampton, SO17 1BJ, UK
[3]Emerging Technology Services, IBM United Kingdom Ltd, Hursley Park, Winchester, SO21 2JN, UK
[4]Department of Computer Science, Rensselaer Polytechnic Institute, Troy, NY 12180, USA.

Corresponding author: ps02v@ecs.soton.ac.uk

Abstract

In thinking about the transformative potential of network technologies with respect to human cognition, it is common to see network resources as playing a largely assistive or augmentative role. In this paper we propose a somewhat more radical vision. We suggest that the informational and technological elements of a network system can, at times, constitute part of the material supervenience base for a human agent's mental states and processes. This thesis (called the thesis of network-enabled cognition) draws its inspiration from the notion of the extended mind that has been propounded in the philosophical and cognitive science literature. Our basic claim is that network systems can do more than just augment cognition; they can also constitute part of the physical machinery that makes mind and cognition mechanistically possible. In evaluating this hypothesis, we identify a number of issues that seem to undermine the extent to which contemporary network systems, most notably the World Wide Web, can legitimately feature as part of an environmentally-extended cognitive system. Specific problems include the reliability and resilience of network-enabled devices, the accessibility of online information content, and the extent to which network-derived information is treated in the same way as information retrieved from biological memory. We argue that these apparent shortfalls do not necessarily merit the wholesale rejection of the network-enabled cognition thesis; rather, they point to the limits of the current state-of-the-art and identify the targets of many ongoing research initiatives in the network and information sciences. In addition to highlighting the importance of current research and technology development efforts, the thesis of network-enabled cognition also suggests a number of areas for future research. These include the formation and maintenance of online trust relationships, the subjective assessment of information credibility and the long-term impact of network access on human psychological and cognitive functioning. The nascent discipline of web science is, we suggest, suitably placed to begin an exploration of these issues.

1 Introduction

The advent of large-scale information networks and portable computing devices has exerted a major influence on many spheres of human activity. Within the space of a generation, the internet has largely transformed the way we work and communicate. But could this new network-enabled world change more than the general structure of our everyday activities? Could its transformative potential extend beyond the realms of the economic and the social to encompass the realms of the psychological and the cognitive? Is the internet capable of fundamentally transfiguring the space of human thought and reason, and, if so, what is the nature of this influence? In this paper, we begin to explore these questions by examining the extent to which notions of environmentally-extended cognition apply to the world of network-enabled information appliances and network-accessible resources (NARs[1]). We argue that, under certain circumstances, the informational and technological elements of a large-scale information network, such as the World Wide Web (WWW), can form part of the material substrate that undergirds a human agent's mental states and processes. Following Clark and Chalmers (1998), we argue that the causally-active physical vehicles of mind and cognition are not always housed entirely within the human body or the human brain; rather, they can (at least sometimes) extend beyond the bounds of skin and skull to encompass elements of the external technological environment.

At the heart of our claim is the (rather radical) proposal that human mental states and processes can inhere in extended, extra-organismic systems that incorporate aspects of the external, technological environment. Claims about the distributed or environmentally-extended nature of human cognition are commonplace in the philosophical and cognitive scientific literature (Clark, 1997, 2003; Clark & Chalmers, 1998; Dennett, 1996; Haugeland, 1998; Hollan et al., 2000; Hutchins, 1995; Kirsh, 1996, 2006; Norman, 1993; Rowlands, 2003; Wilson, 1994; Wilson & Clark, 2006), but the idea that most strongly characterizes our specific claim is that human mental states and processes are, in part, *constituted by* aspects of the external environment. This is the central idea behind the extended mind hypothesis (Clark & Chalmers, 1998). If the hypothesis is true then the opportunities for technology-mediated modes of cognitive transformation seem significant, for inasmuch as our mental states and processes can be said to supervene on elements of the local environment then the manipulation of that environment can be seen as more-or-less akin to the direct manipulation of elements of the human mind. According to the extended mind hypothesis, changes to the environment can, under certain circumstances, exert a profound influence on our long-term store of thoughts, beliefs, memories and knowledge.

The extended mind hypothesis was initially formulated with respect to a number of technologically low-grade examples, mostly involving pen and paper artefacts. Our claim in the current paper is that the extended mind hypothesis can also apply to the realm of network-enabled devices (NEDs) and network-accessible information content (NAIC). Our specific claim is captured in the form of what we will refer to as the thesis of network-enabled cognition:

> Thesis of Network-Enabled Cognition: The technological and informational elements of a large-scale information network can, under certain circumstances, constitute part of the material supervenience base for an agent's mental states and processes.

The thesis of network-enabled cognition implies that the extended mind hypothesis is basically a true account of human cognition and that it can be used to both understand the cognitive impact of network technologies on the human mind as well as guide technology development efforts in support of cognitive and epistemic augmentation. Its validity depends, of course, on the extent to which the extended mind hypothesis actually applies to situations in which both NEDs (e.g. Smartphones, palm pilots, laptops, etc.) and NAIC are exploited by human agents as part of their daily cognitive endeavours. And herein lies the key challenge to the thesis of network-enabled cognition, because it is by no means clear that extended mind accounts do apply in the case of network environments. One concern is that there are fundamental differences between the kind of cognitive/epistemic artefacts used to substantiate accounts of the extended mind (e.g. maps, pens, papers, slide rules, etc.) and the technological (i.e. NED) and informational (i.e. NAIC) components of

[1] A network-accessible resource (NAR) is any resource that can be accessed by a network. It includes the notion of information content, web services, and web applications.

a large-scale information network. Such differences include the availability, accessibility and quality of the information components, as well as the reliability and resilience of the technological elements. There are also concerns about the fact that NARs exist within a common and contested space, one that is vulnerable to sabotage and subversion by cognitively hostile agents (Sterelny, 2004). Even the main protagonists in the extended mind debate seem to eschew (or at least downplay) the possibility that internet resources, such as Google, can fulfil the conditions required for basic forms of cognitive extension (Clark, 2006; Clark & Chalmers, 1998).

The key challenge then is to assess whether the thesis of network-enabled cognition can be defended against claims that network technologies fail to meet the basic criteria for cognitive extension. This is our primary objective in the current paper. We first provide an overview of the extended mind hypothesis and outline some of the criticisms that have been levied against it (see Section 2). We then review the criteria used to identify cases of *genuine* cognitive extension and assess the specific problems that arise when considering the application of these criteria to the various components of a network system (see Section 3). As we will see, most of the concerns about network technologies, from the perspective of the extended mind, centre on the issue of trust. In order to address this specific worry we propose a solution strategy that emphasizes the utilization of computational trust mechanisms and automatic information vetting procedures.

Section 4 explores the idea that the active manipulation of external, extra-organismic resources is an essential element of environmentally-extended cognition. We suggest that in order for network system components to be constitutively-relevant to human cognitive processing they need to provide greater opportunities for information manipulation and personalization. We additionally explore the notion that initiatives such as Web 2.0, which feature greater participatory opportunities than the conventional Web, can be harnessed to provide a greater degree of active restructuring and content personalization so as to meliorate episodes of network-mediated cognitive processing.

Finally, Section 5 summarizes our main conclusions and contributions to the scientific and philosophical literature. Our main conclusion is that network technologies are, at least in principle, capable of meeting the criteria for cognitive extension, and, as such, they are apt for forms of bio-technological hybridization that promise (or threaten) to radically transform the nature of our cognitive profiles. Such mergers, we suggest, introduce a range of issues for future research within the empirical framework of the network and information sciences. In particular, we argue that many of the issues concerning the impact of large-scale networks on human psycho-cognitive functioning are topics apt for exploration within the nascent discipline of web science (Berners-Lee et al., 2006b; Hendler et al., 2008).

2 The Extended Mind

In the attempt to understand human cognition, cognitive science has tended to focus on the brain as the mechanistic substrate of mental phenomena. Cognitive processes, as well as the familiar elements of mentalistic discourse – the mental states we use to explain and predict human behaviour – seem to be firmly located within the traditional biological boundaries of skin and skull. In contrast to this view, Clark and Chalmers (1998) argue that mental states and processes are not necessarily tied to the biological brain; they can instead extend into the external environment to incorporate a variety of external technological props, aids and artefacts. This thesis is referred to as the thesis of extended cognition, or, more commonly, the extended mind. Extended mind accounts typically focus on the way in which much real-world cognitive processing seems to depend on the exploitation of both bodily contingencies and aspects of the external environment (Clark, 1997). Thus, in solving long multiplication problems human agents typically resort to using pen and paper to store intermediate solutions as well as structure the overall sequence of information processing steps (Rumelhart et al., 1986). Similarly, in writing an academic paper the human agent is often engaged in a complex sequence of iterated interactions with a variety of external resources (word processors, books, post-it notes, marginalia, etc.) that, in conjunction with the operation of the human brain, serve to progressively restructure and refine a preliminary set of initial thoughts into a finished article. What these and other examples (Clark, 1997; Hutchins, 1995; Kirsh, 1995; Kirsh & Maglio, 1994) highlight is the importance of external tools, props and artefacts in shaping, augmenting and guiding cognitive processes. Indeed, the extended mind hypothesis argues that such external resources are not only supportive of cognition; they also form part of the very machinery that makes mind mechanistically possible.

In developing the extended mind thesis, Clark and Chalmers (1998) introduce the notion of the parity principle. The parity principle states that:

> *"…if, as we confront some task, a part of the world functions as a process which, were it to go on in the head, we would have no hesitation in accepting it is part of the cognitive system, then that part of the world is (for that time) part of the cognitive process." (Clark & Chalmers, 1998)*

It is this notion that enables Clark and Chalmers (1998) to treat the case of pen-and-paper-mediated long multiplication as a genuine case of environmentally-extended mathematical reasoning. Similarly, the use of an external rotational device to manipulate falling zoids in the game of Tetris (see Kirsh & Maglio, 1994) emerges as a genuine instance of environmentally-extended spatial reasoning. The basic idea is that the parity principle should be used to avoid prejudices and biases about the importance of the inner (neural) realm; it is intended, primarily, as a location-agnostic criterion for identifying the causally-active physical vehicles that undergird a particular mental state or process.

Taken as a basic claim about the importance of external artefacts to adaptive behavioural and cognitive success, there is nothing particularly radical about the notion of the extended mind – many theorists have emphasized the importance of the broader social and technological environment in understanding human cognition (e.g. Hollan et al., 2000; Hutchins, 1995; Kirsh, 1996; Suchman, 1987). What makes the extended mind hypothesis somewhat unique in the philosophical and cognitive scientific literature is the position it adopts with regard to mental states, specifically beliefs. According to Clark and Chalmers (1998), external resources can form part of the long-term store of dispositional beliefs and knowledge associated with a human agent, providing that such resources serve to guide sequences of thought and action in a manner that is functionally equivalent to information retrieved from long-term bio-memory. Take, as an example, a portable device that is used to store information about future meeting appointments. Inasmuch as the use of this device fulfils a number of conditions (more on which below) we are justified, so claim Clark and Chalmers (1998), in seeing the device as akin to a basic biological resource, such as memory. In providing information, the portable device serves the same causal role in coordinating action that the same information would do were it to be retrieved from long-term bio-memory, and, as such, the device can be treated as a store of long-term beliefs about the timing and location of future meeting appointments. According to the extended mind hypothesis, the technological accoutrements of the modern age – palm pilots, calculators, mobile phones, etc. – are not simply tools that support our cognitive endeavours; they are also (potentially at least) part of the material fabric that physically realizes human mental states and processes. Inasmuch as this is the case, the transformative potential of network technologies and resources is indeed profound: by developing technologies that can be easily co-opted into the processing loops of an extended cognitive system, engineers are providing opportunities for cognitive extension that promise (or threaten) to transfigure our traditional notions of cognitive and epistemic (i.e. knowledge-guided) competence.

The extended mind hypothesis sounds controversial, and it is indeed the focus of an ongoing debate within the philosophical and cognitive science communities (Adams & Aizawa, 2001, 2007; Clark, 2005; Clark, 2006; Clark, 2007; Menary, 2006; Rupert, 2004). One reason why the extended mind hypothesis may initially seem unpalatable is that many equate cognition with consciousness, and it seems intuitively wrong to insist that consciousness could supervene on a motley collection of external artefacts. Harnad (2005), for example, disputes the claim that cognition could ever extend beyond the bounds of the body. He claims that "Cognition is thinking; it feels like something to think, and only those who can feel can think" (Harnad, 2005). Clark and Chalmers (1998) accept that the content of occurrent mental states may supervene on processes inside in the brain, but they reject the notion that feelings are important from the perspective of extended cognition:

> *"…not every cognitive process, at least on standard usage, is a conscious process. It is widely accepted that all sorts of processes beyond the borders of consciousness play a crucial role in cognitive processing: in the retrieval of memories, linguistic processes, and skill acquisition, for example. So the mere fact that external processes are external where consciousness is internal is no reason to deny that those processes are cognitive." (Clark & Chalmers, 1998)*

Another (common) criticism of the extended mind hypothesis concerns the lack of any functional equivalence between a purely internal process (e.g. retrieving information from long-term biological memory) and a process that involves the exploitation of an external resource (e.g. retrieving information from an electronic appointment scheduler)[2]. The argument here is that a process like biological memory involves more than the static retrieval of information items; it is an active, dynamic process that both benefits and suffers from the vagaries of mood, motivation and contextual influences. The case of the portable appointment system doesn't seem to fit this specific profile of performance and processing constraints, and it is not, therefore, equivalent to bio-memory. In responding to this criticism, it is important to recognize that nothing in the extended mind thesis commits us to saying that external resources are *identical* to the processing elements of the bio-cognitive system. Clearly they are not. What is important, however, is that the external resource plays the same functional role as inner resources in guiding and constraining action and decision-making. Inasmuch as the information provided by the appointment scheduler system guides and constrains behaviour in much the same way as it would do were it to be retrieved from long-term memory, it does seem sufficient to regard the external information as part of the human agent's repository of knowledge and beliefs about future appointments (see also Clark, 2006). Of course, the notion of functional role similarity is not necessarily without its problems. In particular, it is important to understand the precise way in which inner and outer resources play the same functional role in guiding action and decision-making – simply saying that they are functionally

[2] Less Carr (personal communication).

equivalent, but not at the process level, seems uncomfortably vague. We suggest that the notion of functional similarity should be grounded in the extent to which the external resources support the expression of behaviours that warrant the ascription of mental states to the agent in question. That is to say, what makes an external resource part of the cognitive machinery of an agent is the fact that it plays a causally-active role in the production of behaviours that are the *legitimate* targets of belief, and, more generally, thought ascription.

Other criticisms of the extended mind thesis concern the location of cognitive and computational control (Butler, 1998), the distinction between intrinsic and derived content (Adams & Aizawa, 2001, 2007), and worries about the vulnerability of external resources to damage and social manipulation (Sterelny, 2004). All of these concerns have been addressed by Clark in a number of recent publications (Clark, 2005; Clark, 2006; Clark, 2007). The specific concern raised by Sterelny (2004) about the vulnerability to active sabotage is addressed in more detail below.

3 Extended Cognition and Network Technologies

One of the biggest worries about the extended mind hypothesis is that it is overly liberal concerning what can and cannot be considered part of the cognitive profile of a human agent. Clearly, not all of the technologies or external resources that we encounter are apt to engage in the kind of bio-technological hybridization envisioned by the extended mind hypothesis. As Clark (1997) argues "There would be little value in an analysis that credited me with knowing all the facts in the Encyclopaedia Britannica just because I paid the monthly instalments and found space for it my garage" (pg. 217). Similarly, it would be foolish to equate my personal body of knowledge as co-extensive with the contents of the internet simply because I have an internet-enabled mobile phone. What, then, are the conditions under which we count a set of external resources as participating in genuine cases of cognitive extension? In answering this question, Clark and Chalmers (1998) propose the following set of criteria:

1. "…the resource must be available and typically invoked" (Clark, 2006). **[Availability Criterion]**
2. "…any information…retrieved from [the non-biological resource must] be more-or-less automatically endorsed. It should not usually be subject to critical scrutiny (unlike the opinions of other people, for example). It should be deemed about as trustworthy as something retrieved clearly from biological memory" (Clark, 2006). **[Trust Criterion]**
3. "…information contained in the resource should be easily accessible as and when required" (Clark, 2006). **[Accessibility Criterion]**

As an example of a system that meets all these criteria, Clark and Chalmers (1998) present the case of a tightly coupled system involving a human agent and a conventional (paper) notebook. The notebook, Clark and Chalmers (1998) argue, represents the kind of portable device that supports the retrieval of facts and memoranda in a way that is similar to internal memory (recall the discussion about functional role similarity in the previous section). As such, the notebook-plus-person system represents a genuine case of cognitive extension.

Do the elements of a network system, i.e. NEDs, NARs and NAIC, fulfil (or at least potentially fulfil) the criteria associated with an extended mind? In answering this question it is important to understand the factors that might motivate a distinction between the case of the notebook-plus-person system and the case of the network-plus-person system. If there are no significant differences between the two systems then there is no real point to argue – a network-plus-person system is just as capable as constituting a genuine case of cognitive extension as the notebook-plus-person system. Unfortunately, however, there are many differences between the two cases. In the case of NAIC, significant differences include: physical location (remote vs. local), extent of information sharing (NAIC is typically shared and used by multiple individuals), provenance (NAIC is, in many cases, created by agents who are not necessarily the same as those who consume the information), information quality (NAIC may be of more variable quality than the notebook-derived information) and degree of personalization (NAIC is seldom personalized). Similarly, there are a host of differences concerning the features of the physical device (for example, notebook versus mobile computing device) used to access information. These include power dependence, resistance to damage, reliability of information provision, ease of use, and so on. What significance do these differences have in terms of the validity of the thesis of network-enabled cognition? The general consensus in the literature seems to be that network systems make poor candidates for genuine cases of cognitive extension. Thus in considering the extent to which the internet fulfils the aforementioned criteria, Clark and Chalmers (1998) argue:

> *"The Internet is likely to fail on multiple counts, unless I am unusually computer-reliant, facile with the technology, and trusting, but information in certain files on my computer may qualify."*[3]

In contrast to this rather pessimistic conclusion, we suggest that near future extensions of systems like the internet could fulfil the criteria for cognitive extension (although we also acknowledge the limitations of extant systems). In subsequent sections we explore the general profile of network usage contexts with respect to the extended mind criteria outlined above.

3.1 Availability & Use

The first of Clark and Chalmer's criteria (1998) concerns the extent to which an external resource[4] is both available and typically used by a human agent. In order for the informational and technological elements of a network system to feature as part of an extended cognitive system they should, according to the availability criterion, be more-or-less constantly available for us. Part of the concern here relates to the portability of the device used to access information – the more portable the device, the more likely it is to feature as part of an extended cognitive system. Clearly, the current state-of-the-art in mobile computing features a number of devices that are about as portable as a typical (paper) notebook (the device used in the original account of the extended mind hypothesis). Well known examples include Personal Digital Assistants (PDAs), Smartphones and ultra-mobile personal computers. In terms of the extended mind hypothesis, therefore, portability does not seem

[3] The difference between local and remote files is presumably important here because of accessibility (can we access the information in a suitable timeframe?), reliability (can we always access the information when we need it?), and trust-related (we are more likely to trust local information than we are information stored in a large-scale, publicly-accessible information repository) concerns.

[4] Clark's notion of a 'resource' here is not necessarily equivalent to our notion of a NAR (a web service or institutional website, for example). It is (we think) roughly analogous to the device used to access online information content (i.e. a NED, such as a Smartphone). In the case of a paper notebook there is a tight coupling between the resource used to access information content and the information content itself – the location of the information content is always the same as the location of the notebook. This is not necessarily the case in network-enabled environments; the physical location of information content need not be the same as the device used to access the information.

to be a major problem. In fact, mobile computing devices are already in widespread use and they have largely replaced notebooks and filofaxes as devices that guide daily action and decision-making.

Another concern that may be raised about network resources is their vulnerability to environmental damage. The paper notebook seems somewhat more resilient here as an information access device simply because it does not rely on a host of sensitive electronic components. Of course, claims about the differential vulnerability of devices to environmental damage are not entirely unproblematic: it is not clear that modern electronic devices are any more or less vulnerable to damage than the typical resources (e.g. pen, paper, slide rules, maps, etc.) mentioned in extended mind accounts. In addition, it is worth pointing out that we do not regard the vulnerability of our own nervous systems – systems that are, in fact, highly vulnerable to the ravages of injury, disease, trauma, fatigue, etc. – as necessarily undermining their status as legitimate parts of a cognitive system. One thing that determines the vulnerability of a device to physical damage is the extent to which it is shielded (actively or otherwise) from harm. We suggest, therefore, that while a certain degree of physical resilience is important, a further factor is the extent to which a device is valued and actively protected by its human user. Devices that play a cognitively and epistemically-potent role in the life of a human agent are likely to elicit the kind of protective actions that extenuates the significance of any physical vulnerabilities[5].

Concerns about power consumption and battery life are also relevant to a consideration of resource availability and resource usage issues. One of the obvious differences between a paper notebook and a typical NED relates to the reliance of the latter on a source of power – cases where patterns of resource usage are dictated by power considerations are, unfortunately, all too familiar. One option here is to accept the current shortcomings of NEDs with respect to power constraints, but to insist that such shortcomings represent the active targets of ongoing technology development programmes. A number of recent developments seem of interest and value here. These include the development of energy efficient devices, the exploitation of novel power generation mechanisms (for example, kinetic energy solutions[6]), and the development of novel energy transfer and power sharing solutions (e.g. wireless power transfer solutions (Kurs et al., 2007)).

In summary, when comparing NEDs with non-electronic physical devices, it does seem as though there are some cases where NEDs might fail to fulfil the availability criterion. This is despite the fact that the portability of modern mobile computing solutions is roughly equivalent to, if not greater than, the technologically low-grade examples used in accounts of the extended mind. Particular problems include the power constraints of current NEDs and their relatively greater vulnerability to damage when compared to non-electronic devices. These problems do not, however, seem to rule out the possibility that NEDs could contribute to cognitive extension; rather, they seem to reflect the limitations and operating characteristics of a particular class of device, namely that associated with the current state-of-the-art. To our mind, once these problems have been ameliorated by (near)

[5] It should also be pointed out, of course, that network resources are not always more vulnerable to damage and destruction than their non-networked, non-electronic counterparts. The existence of remote back-up capabilities means that networked information is somewhat more resilient than information stored in a physical notebook. The tight coupling between the notebook and its information contents means that if the notebook is lost or stolen then its contents are also lost. This is not the case with NAIC, which is largely immune to localised forms of damage, destruction and theft.

[6] http://www.m2epower.com/

future innovations, there does not seem to be any principled reason why NEDs could not fulfil the conditions of the availability criterion, at least to the same extent as a non-electronic artefact.

3.2 Trust

The second of Clark and Chalmers' (1998) criteria concerns the notion of automatic endorsability or trust. The idea here seems to be that information retrieved from external, non-biological resources should always be trusted to the same extent as information retrieved *clearly* from biological memory. This is an interesting criterion, because it may seem that we can never trust external resources to the same extent as we trust our own internal (neural) resources. Given that most NAIC exists in a shared and largely public space, it seems inevitable that it will be subject to greater critical scrutiny and epistemic vetting simply because the possibility for deception or deliberate manipulation seems greater in an external, shared space than it does for the inner, private realm.

In considering the trust criterion it is important to be clear exactly what is meant by the terms 'trust' and 'automatic endorsability'. Trust has been the focus of considerable attention in the human-computer interaction literature (see Golbeck, 2008), but it is notable that little consensus exists as to what the term actually means (Corritore et al., 2003). Corritore et al (2003) therefore define (online) trust as an expectation that one's vulnerabilities will not be exploited, but it is unclear whether this definition is compatible with the notion of trust used by Clark and Chalmers (1998). The notion of automatic endorsability is similarly vague. It is probably intended to mean that information should be automatically accepted by an agent without critical scrutiny or cross-checking. This mode of operation seems to best characterize the way we treat information retrieved clearly from bio-memory, and the same mode of operation should apply, it seems, to cases where we rely on external information to guide sequences of thought and action.

Although some commentators have questioned whether externally-derived information can ever be treated in the same way as internally-derived information, it is by no means clear that human agents *always* treat externally-derived information in a highly cautious manner. The psychological phenomenon of change blindness, for example, attests to the fact that people are, at times, remarkably vulnerable to certain forms of deception (Simons & Levin, 1997) (our susceptibility to stage magic tricks is surely a further example of this vulnerability (Kuhn et al., in press; Macknik et al., in press)). Such vulnerabilities could not be exploited, it seems, if humans *always* treated externally-derived information with a high degree of critical scrutiny and scepticism.

In terms of the everyday use of network resources, it is unclear to what extent externally- and internally-derived information is treated in the same way. The issue of whether information from network resources is ever deemed as trustworthy as information retrieved from biological memory is a moot point because, to our knowledge, there have been no empirical studies that would shed light on the issue[7]. What does seem likely is that we do not uniformly trust or distrust external resources (including network ones); rather, our trust evaluations are highly variable and are based on a range of factors associated with the reliability, quality and credibility of online information content (Corritore et al., 2003). The key issue (it seems to us) is not whether we treat external information differently to internal information (to which the answer is surely sometimes we do and

[7] It is worth pointing out that just because an external resource is trusted it does not necessarily follow that information from it will be automatically endorsed. An agent could trust an external resource and yet still perceive the information it receives from the resource as, on occasion, unreliable, inaccurate or of poor quality.

sometimes we don't); it is more a case of when, and under what circumstances, we treat internal/external information in different ways. In answering this question it is clearly important to understand the psychological mechanisms that contribute to information reliability and credibility assessments in network environments. It is also important to understand the processes that undergird the formation and maintenance of subjective trust evaluations in online contexts. Such issues are the current focus of research in the human-computer interaction and web science communities, and we expect the results of this research to further our understanding of the kind of conditions under which genuine cases of network-enabled cognitive extension can be established.

One further point of interest with regard to the trust criterion is the work of Kim Sterelny (2004). Sterelny (2004) is one of the main supporters of the claim that externally-derived information is never trusted to the same extent as internally-derived information. According to Sterelny (2004), the greater vulnerability of external resources to sabotage and subversion means that agents need to deploy strategies that safeguard against the possibility of deception and deliberate manipulation whenever they retrieve information from the external environment. This requirement to deploy safeguards is significant, Sterelny (2004) argues, because it results in an increase in cognitive load and undermines the extent to which external resources can be treated in a similar fashion to internal resources.

In considering Sterelny's (2004) claim we suggest that inasmuch as the use of external resources does indeed increase cognitive load[8], the use of computational trust mechanisms and information quality evaluation techniques can be used to reduce this load. In particular, we suggest that much of the cognitive load incurred by agents in network usage contexts is likely to derive from uncertainty about either 1) the trustworthiness of an external resource or 2) the accuracy, reliability and quality of externally-derived information. Both these sources of load can be attenuated, we argue, by the use of computational trust mechanisms and automatic information vetting procedures (e.g. Adler et al., 2007; McGuinness et al., 2006). One strategy for reducing load is to rely on techniques that evaluate online information resources and provide explicit feedback about the quality, reliability and trustworthiness of specific information items. This is the approach adopted by Adler et al (2007) in the case of Wikipedia content. Another strategy is to highlight the information items that support a *user's* subjective assessment of the quality, reliability and credibility of information content. Liu (2004) thus notes that information about academic institutions plays a key role in user's assessments of the credibility of scholarly information content. He suggests that such information should be highlighted in order to support user credibility assessments. Finally, it is clear that robust security mechanisms are needed to protect online information content. One of the concerns of Sterelny (2004) is that information in the public domain is vulnerable to tampering. When such tampering is

[8] It is worth pointing out that the use of external resources does not always seem to entail a correlative increase in cognitive load. In their analysis of the use of an external rotational device in the game of Tetris, for example, Kirsh and Maglio (1994) calculate that the physical rotation of a shape through 90 degrees takes about 100 milliseconds, plus about 200 milliseconds to actually select the rotation button. To achieve the same result by mental rotation alone takes about 1000 milliseconds. This seems to be a clear case where the use of an environmental resource reduces the cognitive burden associated with a problem-solving task. Daniel Dennett (1996) argues that such forms of environmental off-loading constitute a key feature of human cognition: "The primary source [of our greater intelligence]. . . is our habit of off-loading as much as possible of our cognitive tasks into the environment itself – extending our minds (that, is our mental projects and activities) into the surrounding world".

made more or less impossible by security protection schemes, one need not worry that the information stored therein will be contaminated.

Our basic claim, then, is that inasmuch as the use of network resources entails an increase in cognitive load, computational trust mechanisms and automatic information vetting procedures can be used to attenuate[9] the cognitive burden. Inasmuch as such mechanisms facilitate the automatic endorsement (or rejection) of information content, they are, we argue, important elements of the psycho-technological matrix that supports extended cognition in a network-enabled environment.

3.3 Accessibility

The third of Clark and Chalmers' (1998) criteria concerns the notion of accessibility. The idea here is that information should be easily and readily accessible, perhaps no less so than the information contained in biological memory. At the heart of the accessibility criterion is the notion of information access cost. As the costs (e.g. time and cognitive effort) associated with information access increase then it seems less likely that external information will be seen to count towards an individual's long-term store of knowledge and beliefs.

In empirical analyses it has been observed that systematic manipulations of information access cost can influence the kind of cognitive strategy used by human subjects to solve a problem. Gray and Fu (2004), for example, show that the choice between an internal (memory-based) and external (perceptual-motor) strategy tends to be guided by the relative costs and benefits associated with each strategy (in this case measured in terms of time). If the cost of implementing some perceptual-motor routine is less than the cost associated with the retrieval of information from biological memory then information from the external store will be utilized in preference to the internal information. What these results seem to suggest is that the extent to which external information resources are factored into cognitive processing is heavily reliant on a delicate trade-off between the relative costs of accessing internal and external information. In order for external information resources to become reliably incorporated into problem-solving routines, we should expect their access costs to be minimal; the costs, at least, should be no greater than the time and (perhaps effort) spent in accessing internal information[10].

How do network resources fare with respect to the accessibility criterion? At first blush there seems to be a rather alarming number of problems. The usability of the network device, the time taken to download information content, the vulnerability of the network connection to the vagaries of the

---

[9] This reduction in cognitive effort is consistent with the approach proposed by Parsell (2006). He argues that one approach to reducing the increased cognitive load associated with the vetting of external information is to off-load such vetting mechanisms to the external environment. Our advocacy of computational trust mechanisms represents one potential realization of this proposal.

[10] One thing that is unclear here is the timeframe within which external information should be retrieved so as to count towards an agent's long-term store of knowledge and beliefs. Clearly, if an agent takes too long to access task-relevant information then it is doubtful that they can be credited with knowing or believing that information. But what counts as too long in this situation? One option is to fall back on basic biological standards and to insist that information should be retrieved as quickly as information from biological memory. The problem here is that biological recall can sometimes be quite variable. Another option is to rely on standards in the artificial intelligence community. In talking about the SOAR symbol processing system, for example, Rosenbloom et al (1992) refer to the 'cognitive band' – the timeframe between 10 milliseconds to 10 seconds in which most aspects of deliberative thought seem to occur. Such a timeframe is, we argue, likely to be close to the temporal window within which NAIC needs to be made available in order for mental state ascription to have any explanatory virtue.

radio propagation environment, the effort involved in locating information (e.g. formulating a search query) and the difficulty of isolating relevant information once it has been downloaded: all seem to be factors contributing to an increase in information access cost. We do not have space to address all these issues here; suffice to say that (as with the availability and trust criterion) the focus of many research and technology efforts, particularly in the context of the WWW, are geared towards improving user access to online information. Work of particular note here includes the development of natural language question-answering systems (Lopez et al., 2005; Tablan et al., 2008), user-friendly semantic information browsers (schraefel et al., 2005), the use of sub-vocalization techniques to support Web navigation (Jorgensen & Binsted, 2005) and the use of intelligent forward caching and data charging mechanisms to mitigate download delays and the effects of intermittent network connectivity (Cherniack et al., 2001). New technologies in the field of wearable computing are also likely to enhance our access to information. Mobile device eyewear systems[11], for example, display information directly to a user's visual field using conventional eyewear equipment (e.g. spectacles). Some of the applications envisioned for this new technology include location-aware social network services, real-world visual overlays for environment navigation, battlefield situation awareness displays, and immersive virtual reality systems for education and entertainment. Such systems tend to reduce the cost of information access, and, we argue, they introduce new ways in which NAIC can be co-opted into the information processing loops of cognitively-extended agents.

Notwithstanding the introduction of these emerging technologies, many would argue (ourselves included) that the real source of the information access cost problem (at least in the context of the WWW) lies in the way in which information content is embedded in particular source documents (e.g. web pages) typically using the medium of natural language. This document-centric mode of information encoding is highly impractical for the kind of information access that is likely to be required in the cases of network-extended minds. For what such minds require, we argue, are flexible modes of data integration, aggregation and presentation, in conjunction with an ability to gear information retrieval operations to suit the task-specific needs and requirements of particular problem-solving contexts. Think about the problem of accessing factual information from a web-accessible resource, such as Wikipedia. Even if the delays associated with document retrieval (downloading) and presentation are resolved, the user is still confronted with the onerous task of surveying the document for relevant information content. In most cases this requires the user to scroll through the web page and process large amounts of largely irrelevant content in order to identify the small amount of information that is actually needed. This is a very inefficient means of information access. Even if the user tries to isolate specific information items for use on multiple occasions they cannot do this without reliably fixing the physical location of the information (perhaps by copying the required information to a local resource[12]).

One solution to this problem of resource-centric information embedding is to adopt the kind of approach to data modelling advocated by the Semantic Web community (Berners-Lee et al., 2001; Shadbolt et al., 2006). In this case, domain ontologies are used to represent information in a manner that is largely independent of specific presentational formats or usage contexts. Commenting on the

[11] http://www.microvision.com/wearable_displays/index.html

[12] Links to sections within the page will not work because Wikipedia, like most Web 2.0 applications, features dynamic content, and the physical location of specific information items is liable to change across multiple usage contexts.

relationship of the Semantic Web to the conventional Web, Berners-Lee et al (2006b) write "The SW [Semantic Web] tries to get people to make their data available to others, and to add links to make them accessible by link following. So the vision of the SW is as an extension of Web principles from documents to data" (pg. 18). This shift of emphasis (from linked documents to linked data) is, we suggest, a necessary step in order to establish the kind of selective data integration, aggregation, and filtering that facilitates cost-effective modes of information access and undergirds cognitively potent forms bio-technological integration.

3.4 Summary

The extended mind hypothesis features three criteria that can be used to judge whether an external resource constitutes part of the cognitive profile of an intelligent system. We have suggested that there is no principled reason why the technological and informational components of a network system should not fulfil these criteria. This does not mean, however, that genuine cases of cognitive hybridization are apparent in contemporary usage contexts. In many cases, current network technology seems to fall short of the conditions deemed necessary for cognitive extension. Particular worries concern the power characteristics of mobile devices, the resilience of wireless networks to environmental challenges and the accessibility of information content vis-à-vis flexible modes of information integration, aggregation and presentation. Interestingly, many (if not all) of these shortcomings are the targets of current technology design and development efforts, and this highlights a deep complementarity between the requirements for cognitive extension and the general trend of network technology evolution. Mobile devices are thus beginning to exploit new forms of energy conservation and power generation (Kurs et al., 2007), while initiatives such as the Semantic Web (Berners-Lee et al., 2001; Shadbolt et al., 2006) promise to deliver new approaches to data modelling, ones that liberate data from existing resources and make it available for cognitively potent forms of user interaction.

4 Network-Mediated Cognitive Processing

As alluded to in Section 2, the notion of the extended mind applies to both mental states and cognitive processes. The kind of role that many network systems are likely to play, at least in their current form, is as an environmentally-extended store of long-term knowledge and dispositional beliefs. It is entirely plausible, we suggest, to see elements of a network as effectively expanding the epistemic or belief profile of a human agent (i.e. the kind of knowledge and beliefs the agent is presumed to possess), but it is much harder to see how (extant) network systems could form an intrinsic part of the cognitive processing profile of such an agent. One problem is that NARs, particularly those found on the conventional Web, seem to lack the kind of interactive capabilities that would make them sure-fire candidates for incorporation into extended cognitive systems. Thus, when we confront examples of environmentally-extended cognitive processing, whether it be in the area of memory, perception or reasoning (see Rowlands, 2003), what we seem to encounter are cases in which brain, body and world form a closely-coupled system, one in which aspects of the external environment are manipulated so as to influence subsequent internal processing. The case of Scrabble is a case in point (Kirsh, 1995). What we see in the case of Scrabble is an iterative sequence of body-world interactions in which aspects of the external environment (the Scrabble tiles) are manipulated so as to better exploit the computational profile of the human brain, i.e. its capacity for rapid pattern-recognition and pattern-completion. The problem of word search in this case seems to have been partially transformed from an exhaustive search of the space of cryptarithmetic solutions into something more akin to perceptual pattern recognition and completion: candidate or partial word solutions that are encountered in the perceptual input serve to guide tile manipulation actions in ways that facilitate the further derivation of candidate solutions. Kirsh and Maglio (1994) argue that such environmental restructuring is a common feature of human problem-solving. They argue for a distinction between two types of action: pragmatic and epistemic action. The goal of pragmatic action is to alter the world in such a way as to bring one physically closer to a goal (e.g. one must boil water before making a cup of tea). The goal of epistemic action, in contrast, is to adapt or modify the environment in ways that better suit the problem-solving profile of the agent in question. It is a type of action that is executed in order to make the external environment more *congenial* to cognitive processing (Kirsh, 1996).

The question is to what extent do network resources support the kind of interactive coupling seen in cases of environmentally-extended cognitive processing? What seems relatively clear is that many extant NARs, particularly those associated with the conventional Web, do not provide the same kinds of opportunities for user interaction and manipulation as do many real-world artefacts – a user cannot interact with a standard web page in quite the same way as they can with an abacus or a slide rule, and this, we suspect, limits the range of cognitively potent interactions that a conventional web page can support. A related problem here is that many web pages feature fixed information content that cannot be (easily) adapted, modified or restructured to suit the end user's (task-specific) needs and requirements. As Sterelny (2004) argues:

> *"External tools are in the public domain, and that has implications for the cognitive resources needed to make effective use of those tools. Jointly used epistemic artefacts are often less than optimal for any of their users: they need to be individualised at each use." (Sterelny, 2004)*

While it is certainly true that current opportunities for personalization and customization are somewhat limited in the context of the conventional web, there are a number of efforts underway that promise to change this. The advent of Web 2.0, for example, has increased the extent to which communities can adapt both the content and structure of web-accessible information to suit their information processing needs. Web 2.0 technologies promote greater user participation in the process of editing, uploading and publishing information content, and they therefore enable users to play a greater role in structuring and restructuring Web content so as to better support online problem solving. Nevertheless, it is not entirely clear that the kind of interactive opportunities provided by the current (or even the near-future) Web will necessarily suffice to support the kind of cognitive processing made possible by traditional cognitive artefacts. One possibility is that existing NAIC could be exploited in the context of more interactive environments, although the development of such environments probably depends on more flexible modes of data integration, access, and combination than is currently the case. This suggests another advantage of the 'data liberalization' approach countenanced by the Semantic Web: it makes information more amenable to the kinds of multi-functional, multi-purpose applications that will be needed to support robust forms of user-centred information manipulation and information visualization.

Irrespective of the extent to which interactive technologies support online cognitive processing, an ability to interact with NARs is important for another reason. Inasmuch as interactive technologies increase the participatory nature of large-scale information networks, they may provide a potential solution to the worry that NAIC cannot be trusted to the same extent as internal information (see

Section 3.2). With the advent of more interactive capabilities, users have an ability to create and maintain their own personalized repositories of task-relevant information. In many cases the content of such repositories will have been vetted, checked and approved on previous occasions; it is thus apt for the kind of automatic endorsement required by extended mind accounts (providing, of course, that security mechanisms are sufficiently robust so as to ensure against deliberate tampering and manipulation).

5 Conclusion

The conventional view in cognitive science is that much of the physical machinery of the human mind is located firmly inside the human head. In contrast to this view, the extended mind hypothesis posits that body and world can sometimes form part of the machinery by which mind and cognition are physically realized; body and world hence form part of the local material supervenience base for mental states and processes. The notion of the extended mind is important because it focuses attention on the way in which the facts of material embodiment and environmental embedding (see Haugeland, 1998) may come to shape and scaffold much of what we would otherwise see as the result of purely internal processing. Once we recognize the contribution of the outer, extra-organismic resources, our understanding of the representational and computational profile of the human brain begins to change. For what seems important from this fresh standpoint is the ability of the brain to gear its computational and representational resources to the exploitation (and creation) of external resources that complement, enhance and occasionally transform its computational capabilities (Clark, 1997). The extended mind hypothesis also focuses attention on the opportunities for technology-mediated modes of cognitive augmentation and enhancement. Indeed, one of the distinctive features of the human mind may very well lie in its ability to harness the power of external resources so as to turbo-charge its, otherwise limited, bio-cognitive capabilities. "Human capacities for advanced reason," Clark (2001) thus argues, "depend heavily on the effects of a special kind of hybridization [one] in which human brains enter into an increasingly potent cascade of genuinely symbiotic relationships with knowledge-rich artefacts and technologies".

Given that network systems are now an increasingly pervasive aspect of our lives, coupled with the fact that systems such as the internet have largely transformed our access to information and knowledge, it seems reasonable to ask whether the notion of the extended mind applies to the elements of a large-scale, distributed network system, such as the WWW. We claim that there is, in principle, no reason to doubt this claim – the informational and technological elements of a network system *can* form part of the material supervenience base for an agent's mental states and processes. If this claim is true then the transformative potential of networks with respect to human cognition is indeed profound: the development of new forms of network-enabled capability can be seen as effectively augmenting and extending our basic cognitive and epistemic profiles.

In order to make the thesis of network-enabled cognition stick, we need to evaluate network systems with respect to a number of criteria (see Section 3). In particular, if NAIC is to count as part of our long-term store of knowledge and dispositional beliefs then it is important that the information should be typically used, sufficiently accessible to guide real-time action and decision-making, and more-or-less automatically endorsed. We have seen that if we apply the extended mind criteria to extant network systems, most notably the WWW, we encounter a number of apparent problems. These include (but are obviously not limited to) the extent to which NEDs can be relied on to provide access to online information as and when it is needed, the extent to which online information can be made available for the purposes of guiding real-time thought and action, and the extent to which network resources can be trusted to more-or-less the same degree as personal, biological resources. Rather than seeing these shortcomings as warranting the wholesale rejection of the thesis of network-enabled cognition, we have suggested that these shortcomings in fact reflect the limits of the current state-of-the-art. The shortcomings also identify many of the scientific and technology development goals of ongoing research programmes in the network and information sciences. There is, as such, a significant (and perhaps not coincidental) overlap between the kind of technological advances required to effect potent forms of network-based cognitive hybridization and the goals of many academic and commercial research programmes. Future developments in network systems technology and mobile computing can be expected, we argue, to address shortfalls in the existing accessibility and availability of online resources, and this will constitute a major step forward in terms of our ability to treat network elements as components of an extended cognitive system. Of particular interest are technologies that serve to reduce the effort required to retrieve and process information from network sources. Thus we see highly portable systems, such as Mobile Device Eyewear, establishing a much greater level of perceptual intimacy between human agents and network resources than has heretofore been the case. Similarly, we see advances in sub-vocalization techniques (Jorgensen & Binsted, 2005) and/or neuro-interactive technologies[13] as important steps in enhancing our access to online information and promoting greater levels of cognitive inter-penetration between network resources and the human mind.

Aside from the problems of access and reliability, some of the most significant issues for the thesis of network-enabled cognition seem to relate to the notions of trust and automatic endorsability. Automatic endorsement is deemed to be important because it seems to characterize the way we treat information retrieved (clearly) from bio-memory. If we were to confront a system (either internal or external) in which an agent treated all retrieved information in a highly cautious manner then it seems unlikely that we would see the system as counting as part of the belief or knowledge base of the agent in question. Rather, it is probably the case that we would see the external resource as a separate entity, something akin to a "perceptually-consulted independent oracle" (Wilson & Clark, 2006). As part of our analysis of network-enabled cognition we encountered a number of problems with the notions of trust and automatic endorsability. Firstly, the concepts of trust and automatic endorsability are somewhat vague and ill-defined. The notion of trust, in particular, seems to attract considerable confusion, and there does not appear to be a widely accepted definition of trust in the computer science and human-computer interaction literature (see Corritore et al., 2003). Secondly, there is a dearth of empirical studies that would otherwise shed light on the extent to which NAIC is (or is not) automatically endorsed. It seems clear that people's trust evaluations for various online resources can vary as a function of multiple factors (see Corritore et al., 2003), but this does not, in itself, seem to tell us anything about automatic endorsability: is information from a highly trusted resource subject to automatic endorsement, or is it still treated in a highly cautious fashion? We suspect that what is important here is not so much whether a subject trusts a particular resource so much as the extent to which they deem the information to be credible (or believable). The more credible a human agent perceives information to be, the more likely they are to accept it without question. A final issue that we have found somewhat problematic in relation to trust is the notion that information in the public domain needs to be subject to epistemic vetting and that this necessarily entails an increase in cognitive load (Sterelny, 2004). Inasmuch as the elevated cognitive

[13] http://www.emotiv.com/

load in these cases is attributable to the need to make trust evaluations and/or quality of information assessments, we have argued that technological solutions might be used to ameliorate the cognitive burden imposed on end-user agents (see Section 3.2). What is not clear, however, is whether the use of external resources necessarily entails an increase in cognitive load. Many cases of cognitive off-loading seem, to our mind, to capitalize on the ability of environmental resources to act as stand-ins or surrogates for internal processing. Similarly, some cases of extended cognition seem to involve *less* effort than their internal counterparts (Kirsh & Maglio, 1994). Inasmuch as cognitive load is equated with cognitive effort, it does not seem to us that all cases of cognitive or mental extension necessarily entail an increase in cognitive load. And, even if they do, the use of emerging technologies to provide explicit representations of trust and information quality would seem to constitute an effective method of minimizing this burden. Again, what we seem to confront is a situation where any shortfalls in the capacity of networks to fulfil the criteria for cognition extension turn out to be the targets of ongoing research and technology development in the disciplines of the network and information sciences.

One concern that may be raised here is that we have erroneously conflated the notion of network-based information retrieval with genuine cases of (biologically-based) mnemonic functioning. Even if all the criteria for cognitive extension were to be fulfilled (i.e. the availability, trust and accessibility criteria were fully satisfied), it would be simply false, the argument goes, to regard cases of network-mediated information access as constituting genuine forms of extended memory. And the reason for this is that whereas the information we retrieve from bio-memory will always have been perceived or experienced at the same point in the past (during encoding, for example), the same is not true of much of the information content of a large-scale network environment like the WWW. Most of the information content in this sort of environment will not have been experienced first-hand by the end-user agent. And, as such, it does not appear that we are dealing with a genuine case of extended memory in many cases where such networks are used to guide sequences of thought and action. This is, we think, an important criticism, but it only really applies to situations where we want to talk of networks as directly supporting mnemonic functions. While there are initiatives, like the Memories for Life[14] research programme (O'Hara et al., 2006), that do attempt to explore the contribution of network technologies to human memory, the specific claim we are making about network resources in the context of this paper does not depend on whether such resources constitute a form of extended memory or not. What we are claiming is that NAIC can, under the right circumstances, constitute part of the repository of long-term knowledge and dispositional beliefs associated with a human agent. This claim is neutral with respect to the mechanism by which knowledge and beliefs are actually acquired. In some cases we might want to talk about an agent as acquiring beliefs by virtue of their direct interaction with the world (and their subsequent storage of these beliefs in an online information repository); in other cases, beliefs may be acquired without prior experience – they may be acquired, for example, by the efforts of others to enrich and refine online bodies of communal knowledge[15].

In contrast to the case of mental states, we suggest that contemporary network systems, such as the WWW, are largely ill-suited to supporting environmentally-extended cases of cognitive processing.

[14] http://www.memoriesforlife.org/

[15] Some may see this as a rather radical claim, but the claim that we can progressively enrich an agent's body of domain-relevant knowledge without requiring that agent to experience the knowledge first-hand is largely compatible with much of the work in the knowledge engineering and intelligent agent communities.

We suggest that in the case of such processing it is important that an agent should be able to interact with and manipulate external resources in a highly flexible and largely unconstrained fashion. In particular, an agent should be able to adaptively create and manipulate external resources in such a way that the kind of problem confronting the biological brain is often profoundly different from that confronting the larger cognitive system (involving the brain, body and technological environment)[16]. Initial versions of the WWW were dominated by largely static resources – web pages containing fixed content with limited opportunities for (reconstructive) interaction and manipulation. With the advent of Web 2.0 capabilities we are beginning to witness a gradual move to more participatory modes of information access, ones in which users are encouraged to contribute novel information and enrich existing content using collaborative tagging and annotation schemes (see Steels, 2006). In all likelihood, this move towards greater interactivity on the web is likely to continue, and, as it does so, the opportunities for more cognitively-potent forms of web-based interaction are likely to increase.

In the course of both propounding and evaluating the thesis of network-enabled cognition, our analysis has highlighted a number of areas for future research. Clearly, the factors contributing to subjective information reliability and credibility assessments are important focus areas for research, as are the factors associated with the emergence, maintenance and dissolution of trust relationships. Another area for research relates to the effects of network access on psycho-cognitive functioning. Inasmuch as our mental states and processes do, at times, supervene on the informational and technological elements of a network system then we need to understand (and predict) exactly how our cognitive profile is altered by these basic bio-technological mergers. In particular we need to understand whether the mergers are always positive or even benign (see Carr, 2008). The tendency is to view new technologies in a largely positive light (see Clark, 2003), but as Norman (1993) points out, human agents may come to rely on established technologies in largely unanticipated ways; the sudden replacement of these technologies with newer and seemingly more improved technologies can sometimes prove deleterious. Finally, in exploring the transformative potential of networks we should not ignore the potential effect of network-enabled capabilities on our core notions of who and what we are. If we do indeed stand on the brink of a technological era in which we have easy and reliable access to global knowledge, how will this affect our subjective impressions of our epistemic and cognitive competence? Will it one day make *subjective*, not just scientific, sense to talk about our own personal body of knowledge and beliefs as roughly co-extensive with the contents of a future WWW? And, if so, what effect does this have on our traditional notions of cognitive agency and the self? There are many questions, but, as yet, few answers. Perhaps the real value of the thesis we have proposed is that it focuses attention on some of the problems and opportunities that confront the network sciences and, in particular, the nascent sciences of the Web (Berners-Lee et al., 2006a; Hendler et al., 2008). For the thesis of network-enabled cognition is really a thesis about how the future evolution of our contemporary large-scale information networks may one day come to shape, influence and constrain the basic set of cognitive capabilities we associate with the biotechnologically hybrid mind.

---

[16] This is one way in which external resources exert a transformative influence on human problem-solving: they enable the problem confronting the biological agent to be transformed, whilst preserving the general flavour of bio-computational processing. When we talk about the transformative potential of network systems we should bear in mind that it is not necessarily the basic set of agent-side bio-cognitive capabilities that is being transformed; it is cognitive potential of a system comprising the effective merger of both biological and technological resources.

Acknowledgement

Research was sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence and was accomplished under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the author(s) and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

References

Adams, F., & Aizawa, K. (2001) The bounds of cognition. *Philosophical Psychology*, 14(1), 43-64.

Adams, F., & Aizawa, K. (2007) *The Bounds of Cognition*. Blackwell, Oxford, UK.

Adler, B., Benterou, J., Chatterjee, K., de Alfaro, L., Pye, I., & Raman, V. (2007) *Assigning Trust to Wikipedia Content.* School of Engineering, University of California, Santa Cruz, California, USA.

Berners-Lee, T., Hall, W., Hendler, J., Shadbolt, N., & Weitzner, D. J. (2006a) Creating a Science of the Web. *Science*, 313(5788), 769.

Berners-Lee, T., Hall, W., Hendler, J. A., O'Hara, K., Shadbolt, N., & Weitzner, D. J. (2006b) A framework for web science. *Foundations and Trends in Web Science*, 1(1), 1-130.

Berners-Lee, T., Hendler, J., & Lassila, O. (2001) The Semantic Web. *Scientific American*, 284(4), 34-43.

Butler, K. (1998) *Internal Affairs Making Room for Psychosemantic Internalism*. Kluwer Academic Publishers, Dordrecht, The Netherlands.

Carr, N. (2008) Is Google Making Us Stupid? *The Atlantic*, 302(1), 56-63.

Cherniack, M., Franklin, M. J., & Zdonik, S. (2001) Expressing user profiles for data recharging. *IEEE Personal Communications: Special Issue on Pervasive Computing*, 8(4), 32-38.

Clark, A. (1997) *Being There: Putting Brain, Body and World Together Again*. MIT Press, Cambridge, Massachusetts.

Clark, A. (2001) Reasons, Robots and the Extended Mind. *Mind & Language*, 16(2), 121-145.

Clark, A. (2003) *Natural-Born Cyborgs: Minds, Technologies and the Future of Human Intelligence*. Oxford University Press, Oxford, UK.

Clark, A. (2005) Intrinsic content, active memory and the extended mind. *Analysis*, 65(285), 1-11.

Clark, A. (2006) Memento's revenge: the extended mind, extended. In R. Menary (Ed.), *The Extended Mind*. Ashgate, Aldershot.

Clark, A. (2007) Curing Cognitive Hiccups: A Defense of The Extended Mind. *Journal of Philosophy*, 104(4), 163-192.

Clark, A., & Chalmers, D. (1998) The Extended Mind. *Analysis*, 58(1), 7-19.

Corritore, C. L., Kracher, B., & Wiedenbeck, S. (2003) On-line trust: concepts, evolving themes, a model. *International Journal of Human-Computer Studies*, 58(6), 737-758.

Dennett, D. C. (1996) *Kinds of Minds: Toward an Understanding of Consciousness*. Basic Books, New York, USA.

Golbeck, J. (2008) Trust on the world wide web: a survey. *Foundation and Trends in Web Science*, 1(2), 131-197.

Gray, W. D., & Fu, W. T. (2004) Soft constraints in interactive behavior: The case of ignoring perfect knowledge in-the-world for imperfect knowledge in-the-head. *Cognitive Science*, 28(3), 359-382.

Harnad, S. (2005) Distributed Processes, Distributed Cognizers and Collaborative Cognition. *Pragmatics & Cognition*, 13(3), 501-514.

Haugeland, J. (1998) Mind embodied and embedded. In J. Haugeland (Ed.), *Having thought: Essays in the Metaphysics of Mind*. Harvard University Press, Cambridge, Massachusetts, USA.

Hendler, J., Shadbolt, N. R., Hall, W., Berners-Lee, T., & Weitzner, D. (2008) Web Science: An Interdisciplinary Approach to Understanding the Web. *Communications of the ACM*, 51(7), 60-69.

Hollan, J., Hutchins, E., & Kirsh, D. (2000) Distributed cognition: toward a new foundation for human-computer interaction research. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 7(2), 174-196.

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

Jorgensen, C., & Binsted, K. (2005) *Web Browser Control Using EMG Based Sub Vocal Speech Recognition*. 38th Annual Hawaii International Conference on System Sciences (HICSS'05), Big Island, Hawaii.

Kirsh, D. (1995) The intelligent use of space. *Artificial Intelligence*, 73(1-2), 31-68.

Kirsh, D. (1996) Adapting the environment instead of oneself. *Adaptive Behavior*, 4(3/4), 415-452.

Kirsh, D. (2006) Distributed cognition: A methodological note. *Pragmatics & Cognition*, 14(2), 249-262.

Kirsh, D., & Maglio, P. (1994) On distinguishing epistemic from pragmatic action. *Cognitive Science*, 18, 513-549.

Kuhn, G., Amlani, A. A., & Rensink, R. A. (in press) Towards a science of magic. *Trends in Cognitive Science*.

Kurs, A., Karalis, A., Moffatt, R., Joannopoulos, J. D., Fisher, P., & Soljacic, M. (2007) Wireless Power Transfer via Strongly Coupled Magnetic Resonances. *Science*, 317(5834), 83-86.

Liu, Z. (2004) Perceptions of credibility of scholarly information on the web. *Information Processing and Management*, 40(6), 1027-1038.

Lopez, V., Pasin, M., & Motta, E. (2005) *AquaLog: An Ontology-portable Question Answering System for the Semantic Web*. 2nd European Semantic Web Conference (ESWC 2005), Heraklion, Greece.

Macknik, S. L., Randi, J., Robbins, A., Thompson, J., & Martinez-Conde, S. (in press) Attention and awareness in stage magic: turning tricks into research. *Nature Reviews Neuroscience*.

McGuinness, D., Zeng, H., da Silva, P., Ding, L., Narayanan, D., & Bhaowal, M. (2006) *Investigations into Trust for Collaborative Information Repositories: A Wikipedia Case Study*. Workshop on Models of Trust for the Web, Edinburgh, Scotland.

Menary, R. (2006) Attacking the Bounds of Cognition. *Philosophical Psychology*, 19(3), 329-344.

Norman, D. A. (1993) *Things That Makes Us Smart*. Addison-Wesley, Reading, Massachusetts, USA.

O'Hara, K., Morris, R., Shadbolt, N., Hitch, G., Hall, W., & Beagrie, N. (2006) Memories for life: a review of the science and technology. *Journal of The Royal Society Interface*, 3(8), 351-365.

Parsell, M. (2006) The cognitive cost of extending an evolutionary mind into the environment. *Cognitive Processing*, 7(1), 3-10.

Rosenbloom, P., Laird, J., Newell, A., & McCarl, R. (1992) A preliminary analysis of the SOAR architecture as a basis for general intelligence. In D. Kirsh (Ed.), *Foundations of Artificial Intelligence* (pp. 289-326). MIT Press, Cambridge, Massachusetts.

Rowlands, M. (2003) *Externalism: Putting Mind and World Back Together Again*. Acumen Publishing Limited, Chesham, UK.

Rumelhart, D., Smolensky, P., McClelland, J., & Hinton, G. (1986) Schemata and sequential thought processes in PDP models. In D. Rumelhart & J. McClelland (Eds.), *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Volume 2* (pp. 7-58). MIT Press, Cambridge, Massachusetts.

Rupert, R. (2004) Challenges to the Hypothesis of Extended Cognition. *Journal of Philosophy*, 101(8), 389-428.

schraefel, m. c., Smith, D. A., Owens, A., Russell, A., Harris, C., & Wilson, M. L. (2005) *The evolving mSpace platform: leveraging the Semantic Web on the trail of the memex*. Hypertext 2005, Salzburg, Austria.

Shadbolt, N., Hall, W., & Berners-Lee, T. (2006) The semantic web revisited. *IEEE Intelligent Systems*, 21(3), 96-101.

Simons, D. J., & Levin, D. T. (1997) Change blindness. *Trends In Cognitive Sciences*, 1(7), 261-267.

Steels, L. (2006) Collaborative tagging as distributed cognition. *Pragmatics & Cognition*, 14(2), 287-292.

Sterelny, K. (2004) Externalism, Epistemic Artefacts and the Extended Mind. In R. Schantz (Ed.), *The Externalist Challenge: New Studies on Cognition and Intentionality* (pp. 239-254). de Gruyter, Berlin, Germany.

Suchman, L. (1987) *Plans and Situated Actions*. Cambridge University Press, Cambridge.

Tablan, V., Damljanovic, D., & Bontcheva, K. (2008) *A Natural Language Query Interface to Structured Information*. 5th European Semantic Web Conference (ESWC'08), Tenerife, Spain.

Wilson, R. A. (1994) Wide computationalism. *Mind*, 103(411), 351-372.

Wilson, R. A., & Clark, A. (2006) Situated Cognition: Letting Nature Take its Course. In M. Ayded & P. Robbins (Eds.), *Cambridge Handbook of Situated Cognition*. Cambridge University Press, Cambridge, UK.