Cognition and the Web

Paul R. Smart
School of Electronics and Computer Science, University of Southampton, Southampton, SO17 1BJ, UK.

Introduction

In the few decades since its invention, the World Wide Web has exerted a profound influence on practically every sphere of human activity. Online stores have transformed the way we purchase goods, social networking sites have transformed the way we stay in touch with friends, and real-time news feeds have transformed the way we stay abreast of current affairs. For better or worse, it seems, the Web is poised to have a significant influence on the way we live our lives, and perhaps ultimately it will come to influence the social, political and economic forces that determine the way our lives are lived.

The rapid growth and penetration of the Web raises important questions about its effects, not just on our social activities, but also on the nature of our cognitive and epistemic profiles. The Web is a transformative technology, but its transformative influence does not necessarily stop at the social processes that govern our everyday interactions with one another. Many technologies that have transformed society (for example, the clock, the map and systems of writing), have also exerted subtle (and sometimes not so subtle) effects on our cognitive and intellectual capabilities. The invention of the mechanical clock, for example, contributed not only to widespread social and economic change (see Landes, 2000); it also contributed to a profound shift in the way we saw ourselves and the nature of our cognitive abilities. Prior to the invention of reliable and widely accessible time-keeping devices, the nature of our daily activity was structured according to the chronology of the natural world. The sun signalled the start of the working day, while the onset of night signalled a time to sleep. The invention of portable time-keeping devices changed all that. Suddenly, it was possible to organize and synchronize activities in a way that had never been possible before, and on the back of this new capability there emerged a new social and economic era. The influence of the clock, however, was not just limited to the nature of our social interactions and engagements; it also effected a profound shift in our cognitive power and potential. As Clark (2003) argues:

> *"The presence of easily accessible, fairly accurate, and consistently available time-telling resources enabled the individual to factor time constantly and accurately into the very heart of her endeavours and aspirations. This made possible ways of thought, and cultural practices and institutions, which were otherwise precluded by our basic biological nature." (pg. 40)*

Carr (2010) is also impressed by the cognitive and intellectual repercussions of the clock's invention. He writes that:

> *"The mechanical clock changed the way we saw ourselves. And like the map, it changed the way we thought. Once the clock had redefined time as a series of units of equal duration, our minds began to stress the methodical mental work of division and measurement. We began to see, in all things and phenomena, the pieces that composed the whole, and then we began to see the pieces of which*

> *the pieces were made. Our thinking became Aristotelian in its emphasis on discerning abstract patterns behind the visible surfaces of the material world. The clock played a crucial role in propelling us out of the middle ages and into the renaissance and then the enlightenment." (pg. 43)*

As with the invention of clocks, the emergence of the Web is, I think, a highly significant event in our social, economic and cultural history. Like the invention of the clock, the Web is supporting the emergence of new forms of social interaction and engagement, and ultimately these may manifest themselves in more profound forms of social and economic change. The Web also has the potential to change the nature of our cognitive and epistemic profiles. The Web provides new opportunities for interaction and engagement with a global space of information, and, in some cases, such interactive opportunities may contribute to fundamental shifts in the way we see ourselves and the nature of our cognitive capabilities.

The idea that a technology as pervasive and as popular as the Web may be changing our cognitive profile is something that is bound to raise both interest and alarm. Over the past couple of years, considerable attention has been devoted to the issue of whether the Web is affecting our social, emotional and cognitive abilities. For the most part, much of the rhetoric in this debate has been negative. The Web has been seen as exerting a largely pernicious influence on our ability to think, read and remember, apparently undermining our ability to engage in 'linear thought', and encouraging us to adopt highly superficial forms of information processing (Carr, 2008, 2010). But what is the real cognitive impact of the Web, and should we be alarmed if it stands to fundamentally alter the nature of our cognitive and epistemic profiles?

In this paper, I will attempt to provide an overview of what I see as the important issues in the debate regarding the relationship between human cognition and the Web. I will argue that the Web, like many technologies, is poised to transform the nature of our cognitive and epistemic profiles, and for this reason we should take the views of Carr (2008, 2010) and others very seriously. However, I will argue that much of the negative rhetoric in the debate over the Web's influence on human cognition is unwarranted. The Web does stand to influence the power and potential of human cognition, but in order to fully understand this influence we need to factor in the available scientific evidence, the changing nature of our interaction with the Web, and the possibility that many of our everyday cognitive achievements already depend on complex webs of social and technological scaffolding. I will argue that the worries and concerns of Carr and others are largely unfounded (on empirical and philosophical grounds) and that we should not be overly negative in our characterization of the Web's influence on our cognitive capabilities. Firstly, the scientific evidence that would warrant such a conclusion is simply not available, and much more research needs to be undertaken before we can establish a scientifically informed view. Secondly, even if the evidence does suggest the Web is undermining our cognitive abilities, we need to be aware that the Web is constantly evolving. Current modes of interaction with the Web do not necessarily limit (or even limn) the space of interactive opportunities that could be created by future forms of technological innovation. We should, as a consequence, be very wary of blanket statements to the effect that the Web is undermining our cognitive capabilities. Our notions of what the Web is and what it may yet become are constantly changing, and not all forms of the Web are likely to exert the same kind of influence on our cognitive capabilities. Above all, I want to suggest that when we think about the cognitive impact of the Web, we need to take very seriously the possibility that human cognition is grounded on more than the processing operations of the bare biological brain. Many forms of cognitive competence, I will suggest, are based on the delicate coupling of a variety of resources into complex nexuses of informational flow and influence that span the biological, social and technological domains (Clark, 1997, 2008; Smart et al., 2010a). If such claims are true, then the cognitive capabilities of the human agent are not necessarily those of the technologically-decoupled brain, and the information processing capabilities of the biologically-bounded individual always need to be judged relative to the capabilities of the bio-technological systems in which that individual typically participates. From this perspective, we should not judge the Web solely on the basis of its effects on bio-cognitive processing. Instead, we should attempt to understand more about the cognitive capabilities and limitations of the human individual working in concert with Web technologies. Only then will we be able to judge whether the World Wide Web contributes to our cognitive good or ill.

The structure of this paper is organized around five key issues. In 'The Cognitive Impact of the Web', I attempt to examine the effect of the Web on human cognitive functioning. For the most part, this discussion is based on what we might consider as the conventional Web; i.e. the Web of HTML documents that are traditionally accessed using Web browser technology. I firstly review the concerns raised by Carr (2008, 2010) that the Web, and the Internet more generally, is undermining our ability to think, read and remember. I then go on to look at some of the scientific evidence that might support, or fail to support, such a view. In the section entitled 'Changing the Way We Think?', I consider the specific question of whether the Web is changing the way we think. This question resembles that which was recently posed by the Edge Foundation as part of their Annual Question series[1] (the specific question asked by the Edge Foundation was 'HOW IS THE INTERNET CHANGING THE WAY YOU THINK?'). It is not my aim in this paper to provide a specific answer to this question. Rather, I attempt to highlight a number of issues that affect how the question might be addressed. As is indicated by the broad range of responses to the Edge Foundation's question, the issue of whether the Web is changing the way we think is something that depends very much on how we interpret the phrase 'the way we think', and on what we mean by the 'Web'. In the section entitled 'Extending the Mind: Cognitive Extension and the Future Web', I consider the limitations of the current Web in terms of its ability to support the emergence of extended cognitive systems. This section expands on a discussion that has appeared in a number of places concerning the possibility of Web-based forms of cognitive extension and the emergence of 'Web-extended minds' (Smart et al., 2008; Smart et al., 2009; Smart et al., 2010a). In the context of this paper, I seek to identify those features of the Web that present barriers to the emergence of extended cognitive systems, and I attempt to provide a vision of the future Web in which most of these barriers are removed. In the section entitled 'Creativity and Collective Cognition', I focus on the potential effects of the Web in relation to more socially-distributed forms of cognitive processing; in particular, I focus on the way in which the Web may help (or hinder) some of the social processes associated with creativity. Finally, in the section entitled 'Cognition and the Web: An Agenda for Web Science', I attempt to highlight some of the issues for future scientific work in this area. Given that recent scientific effort around the Web has coalesced into a new scientific discipline, namely that of Web Science (Berners-Lee et al., 2006; Hendler et al., 2008; Shadbolt & Berners-Lee, 2008), I suggest that many of the issues relating to cognition and the Web could form part of the emerging Web Science research agenda.

[1] http://www.edge.org/questioncenter.html

The Cognitive Impact of the Web

The question of whether the Web, and the Internet more generally, is affecting aspects of the human mind has, understandably, been the focus of much recent scientific attention and media interest. The Web has rapidly gained in popularity, so whatever its cognitive effects they are likely to be widespread. One of the biggest concerns is that the Web is changing the way we process information, causing us to rapidly process lots of disparate information at a rather superficial level rather than encouraging us to engage in any kind of deep information processing (Carr, 2008, 2010). One of the most vocal supporters of the view that the Internet is changing the way we think is Nicholas Carr. His concern is that our contact with the Web is contributing to a significant, and not necessarily beneficial, change in our cognitive profiles:

> *"As the media theorist Marshall McLuhan pointed out in the 1960s, media are not just passive channels of information. They supply the stuff of thought, but they also shape the process of thought. And what the Net seems to be doing is chipping away at my capacity for concentration and contemplation. My mind now expects to take in information the way the Net distributes it: in a swiftly moving stream of particles. Once I was a scuba diver in the sea of words. Now I zip along the surface like a guy on a jet ski" (Carr, 2008; pg 57).*

In his recent book, "The Shallows: How the Internet is Changing the Way we Think, Read and Remember", Carr (2010) argues that the Web[2] is having a number of effects on human cognition. Instead of enhancing our ability to concentrate, Carr argues that the Web is undermining our capacity for sustained periods of focused attention. And instead of enhancing our ability to think deeply about a topic, the Web is undermining our ability to engage in protracted episodes of what is referred to as 'linear thinking'. The result of these changes is a curtailment and fragmentation of otherwise temporally-protracted episodes of cognitive activity, coupled with the adoption of highly superficial forms of information processing. We have, according to Carr, become far more easily distractible and "thinly spread". In the words of the playwright Richard Foreman, we have become "pancake people"; people whose cognitive resources are now "spread wide and thin, as we connect with that vast network of information access by the mere touch of a button."

The constellation of cognitive changes alluded to by Carr is rather hard to categorize in terms of conventional psycho-cognitive processes; however, in general, the changes are ones that are related to cognitive control functions, comprehension processes and reward-mediated behaviours. In particular, Carr's analysis suggests that our contact with the Web may be producing the following changes:

1. **Shallow and superficial modes of information processing.** Instead of an ability to think deeply about a particular topic or issue, we now process information at a far more superficial level. Whereas before we had the ability to be deeply immersed in a text, now we are limited to more cursory forms of analysis. This kind of deficit can be linked to the 'levels of processing' literature, in which it has been shown that our ability to remember information is linked to the depth at which the information is processed (Craik & Lockhart, 1972). Information that has been subjected to deep, semantic processing can be recalled

[2] Carr's focus is actually the Internet, rather than the Web. Nevertheless, much of his discussion is oriented to the Web.

much better than information that has been processed at a much more superficial level (see Eysenck, 1974).

2. **Heightened distractibility.** Instead of an ability for sustained periods of focused attention, we are now far more susceptible to the distracting influence of task-irrelevant stimuli.
3. **An inability to sustain episodes of cognitive activity.** Instead of an ability to 'stay-on-task', we are now liable to a premature curtailment of temporally-protracted episodes of cognitive activity.
4. **A change in the rewarding impact of conventional printed media.** Instead of finding pleasure in books, we are now easily bored by reading. This intimates at potential differences in the reward and reinforcement processes associated with different media.

Importantly, these changes are assumed by Carr to be relatively enduring. They persist beyond the periods of time that we are actively using the Web, and they thus reflect enduring changes in our cognitive profiles.

Given the nature of the putative cognitive changes, it is clearly important to know something about the validity of Carr's claims. The problem is that many of the claims made by Carr are based largely on anecdotal evidence. There is little direct scientific evidence to support the notion that the Web is currently introducing fundamental changes in our cognitive profile. This is important because even if it could be shown that our cognitive profile is changing, it is by no means clear that the Web itself is responsible for the change. Perhaps few people who make extensive use of the Web would disagree with Carr when he says that we are losing our capacity for deep thinking and sustained concentration. However, it is by no means clear that the Web itself is responsible for these changes. One possibility is that people who make frequent use of the Web are also likely to engage in something called "media multitasking", which involves the concurrent processing of information content from two or more media sources (for example, surfing the Web and processing SMS messages). Media multitasking has been linked to an altered cognitive profile, characterized by the adoption of breadth-based cognitive control strategies (Ophir et al., 2009), and this alteration in cognitive functioning might account for some of the putative cognitive changes alluded to by Carr (certainly some of the deficits described by Carr seem to involve alterations in cognitive control processes). Another issue concerns the effect of age on our cognitive abilities. It is well-known that some aspects of our cognitive profile change with age, and at least some of these changes are consistent with the deficits described by Carr (for example, see Braver & Barch, 2002). If we are trying to understand the reason for an age-related alteration in our own cognitive functioning, then we may search for answers in the external environment rather than ourselves. The Web, in this case, emerges as an obvious scapegoat. It was introduced relatively recently, and it has precipitated profound changes in the way we consume and communicate information. This makes it a very salient and convenient target in our efforts to account for age-related declines in our cognitive capabilities.

Many studies have sought to examine the effect of hypermedia environments on aspects of cognitive function. In one study, for example, Zhu (1999) examined the effect of the number of hyperlinks on learning performance. He discovered that readers were better able to learn from texts that had fewer links (3-7 links per node) rather than a larger number of links (8-12 links per node). This suggests that larger numbers of links may affect a reader's ability to process information content. Other studies have focused their attention on the structural organization of hypermedia environments; for example, the effect of different hyperlink topologies. In general, the results from these studies suggest that, when it comes to the recall of factual content, linear texts (i.e. texts in which there is a linear sequence of nodes and navigation is restricted to 'next' and 'previous' links) contribute to better recall performance compared to other kinds of hypertext (for example, texts in which links are embedded in the text itself and navigation is not restricted to a linear sequence of nodes) (see DeStefano & LeFevre, 2007). Such results may indirectly support some of Carr's claims about the way in which we process information on the Web. If we do resort to a shallow form of information processing – one in which we largely refrain from any kind of deep, semantic processing – then we may expect subjects to recall relatively less when they are confronted with Web-like hypermedia environments[3]. The problem is that it is difficult to know exactly what accounts for the reduced recall performance. It could be that complex hypermedia environments encourage rather superficial forms of information processing, or it could be that other features of the hypermedia environment are interfering with the proper encoding of information content. For example, the visual cues used to highlight the presence of a hyperlink may increase the visual processing demands associated with a hyperlinked resource (especially when these links are embedded in the main body of a text), and this may be contributing to an increase in cognitive load (see DeStefano & LeFevre, 2007).

Another problem with studies that focus on the offline recall of information (i.e. recall in the absence of any kind of interaction with the hypermedia resource) is that they do not necessarily tell us anything about the effect of hypermedia environments on real-world cognitive performance. Clearly, the fact that a particular kind of hypermedia environment (e.g. one featuring a particular kind of link topology) negatively affects the offline recall of presented information *does* tell us something about the effect of Web resources on (bio-)memory processes. It is not clear, however, that these effects really matter when it comes to our everyday use of the Web. Perhaps subjects do not need to remember information when they interact with the Web because the reliable presence of Web-accessible information content effectively functions as a form of external or 'extended' memory (see Smart, 2010). In this case, the kind of performance deficits seen in laboratory situations would not extend to the real world. In addition, it has been shown that subjects will refrain from memorizing information content if the cost of accessing information from the environment is sufficiently low (Gray & Fu, 2004; Waldron et al., 2007). As such, one potential explanation of the results from recall-based studies is that hypertext environments simply reduce the cost of accessing relevant information and therefore discourage the use of internal, biologically-based memorization strategies.

As an alternative to measuring offline abilities (like the ability to recall factual content), we can also measure aspects of a subject's online performance. For example, we can provide the subject with an initial opportunity to browse a hyperlinked information resource, and then ask him or her to locate target information while actively interacting with the resource. In one study of this kind, McDonald and Stevenson (1996) compared navigation performance (measured as the time taken to locate answers to specific questions) in three hypertext conditions. In one condition, subjects were presented with linear text, a linear sequence of textual nodes in which navigation was restricted to successor and precursor nodes. In a second condition, subjects were presented with a strictly hierarchical text, a hyperlinked resource in which the nodes are organized in a hierarchical fashion

[3] This follows from the predictions made by Craik and Lockhart's (1972) 'levels of processing' theory.

(in this condition, hyperlinks support navigation to subordinate or superordinate nodes in the hierarchy). Finally, in a third condition, subjects were presented with nonlinear hypertext. This final condition most closely resembles the state-of-affairs that is encountered on the Web, with hyperlinks embedded in the text supporting navigation to nodes outside of a strict linear or hierarchical sequence. The number of links in each condition was lowest in the linear condition and highest in the nonlinear hypertext condition, with hierarchical hypertext having an intermediate number of links. The results suggested that as the number of links increased, navigation performance decreased. It thus took subjects longer to find target information in the nonlinear hypertext condition than either the linear or hierarchical hypertext conditions.

What the results of studies by McDonald and Stevenson (1996) and Zhu (1999) suggest is that increasing the number of links in a hypertext resource does not seem to promote either the recall of information content or the adoption of more efficient navigational strategies. In fact, the available evidence suggests that the inclusion of lots of embedded hyperlinks may actually undermine a subject's ability to recall information while disconnected from the hypermedia environment. This performance deficit does not seem to be compensated for by an ability to retrieve information more quickly when actively engaged with the environment[4].

In a review of the literature concerning the effects of hypertext features on cognitive processing, DeStefano and LeFevre (2007) propose that the effect of hypertext on cognitive performance is mediated by an increase in cognitive load. They suggest that each time a user encounters a hyperlink, he or she has to make a decision as to whether or not to follow the link, and this increases the cognitive demands on the user. There are a number of ways in which hypermedia environments may increase a user's cognitive load. Firstly, as DeStefano and LeFevre (2007) suggest, users may need to decide whether or not to follow a link whenever it is encountered in a text. Secondly, the mere presence of links, indicated as they are by a variety of visual cues (e.g. a change in font), may impose visual processing demands that are simply not seen in conventional, non-linked textual resources. Thirdly, individual nodes in a hypermedia environment, such as the World Wide Web, typically feature a variety of textual and non-textual elements. More often than not, individual pages on the Web are composed of a confusing mish-mash of multi-media components, many of which are semantically unrelated to the main content of the page. Page animations, adverts, sidebars, inline videos, pop-up boxes, image rollovers, complex menu navigation systems and all manner of decorative gee-gaws present the user with a plethora of opportunities for distraction that simply do not exist in the context of a conventional printed book. Finally, if the user does decide to follow a link, they are presented with the additional task of integrating information between the source and destination node. In the case where the content of the destination node is semantically-unrelated to the content of the source node, the reader experiences a disruption in narrative continuity that, in all likelihood, increases demands on working memory.

One source of evidence that hypermedia environments impose greater cognitive demands on users comes from studies of individual differences in working memory capacity. If hypermedia environments require the use of more processing resources than other information environments,

[4] The findings of McDonald and Stevenson (1996) also raise important questions about the potential effects of link topology on cognitive performance. The links between resources in a hypermedia environment form networks with different structural characteristics, and it is possible that certain kinds of network structure may exert different kinds of cognitive effects.

then individuals that have more of those resources should be relatively less affected than individuals that have fewer resources. Lee and Tedder (2003) tested this hypothesis using students who scored high, medium or low on a test of working memory capacity. The students were exposed to one of three kinds of texts (linear, hierarchical or networked), and the effect of the texts was subsequently measured using the method of offline recall. The results revealed that recall was highest in the linear condition, and that the advantage of the linear text was greatest for students with the lowest working memory scores. The results are therefore consistent with the idea that hypermedia environments can increase working memory demands, thereby causing individuals with low working memory capacity to be especially vulnerable to the cognitive sequelae of hypermedia environments.

The available empirical results on the effect of hypermedia environments therefore suggest that we may process Web-based information content in a way that is different from that seen with conventional printed media. A key problem seems to be the effect of embedded hyperlinks and unrestricted navigation opportunities. When we encounter a resource that features many hyperlinks, it seems that we may experience an increase in cognitive load relative to that which is seen in the case of conventional non-linked resources or printed materials. The result of this increase in cognitive load is a diminished ability to recall and locate relevant target information.

Some of this evidence is compatible with the views of Carr regarding the effects of the Web on our cognitive abilities. Carr (2010) argues, for example, that hyperlinks are instrumental in contributing to a profile of heightened distractibility:

> *"Links don't just point us to related or supplemental works; they propel us towards them. Hyperlinks are designed to grab our attention. Their value as navigational tools is inextricable from the distraction they cause." (pg. 90)*

Carr's main concern, however, is not that the Web is affecting our cognitive abilities during the time in which we are actually using the Web – during the time we are actually online – he is much more concerned about the long-term effects of the Web. In particular, Carr is worried that the Web is undermining our ability to process information obtained from any source, including our ability to process information derived from conventional printed media. What evidence is there for this more enduring effect of the Web on our cognitive functioning?

The evidence is, in fact, limited or non-existent. The only study that really supports the case for enduring technology-mediated changes in cognitive functioning is a study by Ophir et al (2009). Ophir et al (2009) were interested in the effects of something called media multitasking. This is the concurrent consumption of information from more than one media stream; for example, browsing the Web whilst listening to music or editing an online document whilst talking to a colleague on the phone. In order to assess the effect of media multitasking on cognition, Ophir et al (2009) divided subjects into two groups: heavy and light media multitaskers. They then assessed the performance of the two groups on a variety of tasks intended to index cognitive control capabilities. Their results showed that heavy media multitaskers exhibit a number of changes in cognitive control functioning relative to light media multitaskers. In general, heavy media multitaskers are more susceptible to interference from irrelevant stimuli, and they thus seem to possess a tendency for bottom-up, as opposed to top-down, attentional control mechanisms:

> *"The present data suggest that LMMs [low media multitaskers] have a greater tendency for top-down attentional control, and thus they may find it easier to attentionally focus on a single task in the face of distractions. By contrast, HMMs [high media multitaskers] are more likely to respond to stimuli outside the realm of their immediate task, and thus may have a greater tendency for bottom-up attentional control" (Ophir et al., 2009)*

Of course, the method used by Ophir et al (2009) to distinguish between heavy and light media multitaskers is not one that specifically targets the degree of Web usage. In fact, it is entirely possible that none of the heavy media multitaskers were Web users. The results are interesting, however, because they suggest that the heavy use of contemporary communications and media technology can (potentially[5]) induce relatively enduring changes in information processing; i.e. changes that extend beyond a particular usage context. Given that cognitive control is implicated in our ability to appropriately allocate our attention to a task, maintain information in working memory, and control our responses to task-relevant (as well as task-irrelevant) information, it seems that Carr might be onto something when he suggests that some of us are suffering from a form of technology-induced attention deficit disorder. What is not clear, however, is whether these changes stem from our specific use of the Web, or whether they simply reflect a cognitive response to the multimedia medley of the modern world[6].

If our exposure to the Web is contributing to changes in cognitive control functions, then such changes could be revealed by studies examining the neuroanatomical structures implicated in such functions. The brain structure that has been most commonly implicated in cognitive control (or executive functioning) is the prefrontal cortex (e.g. Miller & Cohen, 2001; Roberts et al., 1998). At least one study has indicated that the performance of a common Web-based activity, namely Web searching, is associated with changes in this region (Small et al., 2009). This study used the technique of functional magnetic resonance imaging (fMRI) to compare the brains of subjects who were either experienced or inexperienced Web users. Measurements of brain activity were taken while subjects were engaged in a simulated version of a Web searching task and while they were engaged in reading a conventional text. The results revealed a significant difference in baseline activity between the brains of experienced and inexperienced Web users while they were engaged in the Web searching task but not while they were engaged in the reading task. In particular, experienced Web users showed a much greater activation of frontal cortical structures, relative to the inexperienced users. Interestingly, this difference was eliminated in a subsequent scan after the inexperienced subjects had been asked to engage in Web searching behaviour for a limited period of time (see

---

[5] One alternative reading of the results of Ophir et al (2009) is that individuals with limited cognitive control functions may simply be more inclined, for whatever reason, to become heavy media multitaskers. As such, it is not clear that the concurrent processing of multiple media streams is the thing that actually causes the shift to breadth-based cognitive control strategies.

[6] Inasmuch as it is the exposure to multiple, distinct streams of information content that is actually causing the changes observed by Ophir et al (2009), it may be that our exposure to the Web is, by itself, sufficient to cause changes in executive functioning. This is because many Web pages feature multiple types of media (video, images, text, audio), and, at least occasionally, these media streams communicate semantically-distinct information. A news site, for example, may display live video, scrolling stock quotes, live sports scores and adverts, while the average Facebook page is a veritable smorgasbord of semantically-distinct information content. In our efforts to deal with these multiple media streams, we may be engaging in a form of Web-based media multitasking that is not dependent on the use of distinct media devices.

Small & Vorgan, 2008). After just five days of Web searching experience, the brains of inexperienced Web users resembled those of their more experienced counterparts (see Small & Vorgan, 2008).

In the absence of controlled behavioural studies, it is difficult to know whether the results of Small et al (2009) tell us anything significant about the effects of the Web on cognitive functioning. While the frontal cortex has been implicated in cognitive control functions, it is difficult to know exactly what an increase in its activity actually represents. In general, higher levels of prefrontal cortical activation are seen on tasks typically requiring executive functioning (e.g. Koechlin et al., 2003; Newman et al., 2003), and so it is not easy to argue that a profile of heightened activity necessarily reflects a deficit in cognitive control. Small and Vorgan (2008) recognize this possibility. They suggest that:

> *"...Googling or other technological experiences does sharpen some cognitive abilities. We can learn to react more quickly to visual stimuli and improve many forms of attention, particularly the ability to notice images in our peripheral vision. We develop a better ability to sift through large amounts of information rapidly and decide what is important and what isn't..." (pg. 49)*

The results are also difficult to interpret in terms of the possibility of enduring changes in cognitive functioning that extend beyond the specific use of the Web. The fact that neither group seemed to show differential activity on the reading control task suggests that the changes in brain function were relatively specific to the Web searching task. If this is the case, the results of Small et al (2009) may reflect the adaptation of the brain to some of the task demands associated with the retrieval of online content.

To summarise, the available scientific evidence suggests that certain features of the Web may affect cognitive processing. In particular, the presence of hyperlinks, and the navigational capabilities they support, may interfere with our ability to recall information and develop a deep understanding of information content. However, there is no direct evidence to support the idea that the Web is contributing to long-lasting changes in cognitive function. There is some evidence to suggest that our attempt to process information from multiple media streams may be associated with changes in executive functioning (Ophir et al., 2009). However, such evidence is not necessarily the result of Web use, and there is insufficient evidence to discern the direction of causality. The neuroimaging studies of Small et al (2009) suggest that frequent Web use may be linked to changes in brain function; however, the cognitive significance of such changes is currently unclear. In addition, the similarity in brain activity profiles of both experienced and inexperienced Web users on a reading control task is not consistent with Carr's claim that the Web is contributing to profound changes in general reading processes.

Changing the Way We Think?

The question of whether the Web is producing cognitive change is a question that has attracted considerable recent media attention and intellectual debate. In addition to a continuing commentary in the British media[7], the cognitive implications of Internet access were the focus of the 2010 Edge Foundation Annual Question. In response to the question "HOW IS THE INTERNET

[7] See, for example, http://www.guardian.co.uk/technology/2010/aug/15/internet-brain-neuroscience-debate and http://news.bbc.co.uk/1/hi/technology/7459182.stm.

CHANGING THE WAY YOU THINK?" a variety of commentators, ranging from scientists, philosophers, journalists, historians and technologists, offered their views as to how they thought the Internet was affecting human cognition[8]. Most of these views recognized the potential of the Internet to effect profound shifts in our cognitive abilities.

In this section, I want to highlight a number of issues that I believe are important to consider in answering questions about the cognitive impact of the Web. Ultimately, the nature of the cognitive changes produced by the Web will only be uncovered by empirical research, but in undertaking this research it is important to have a clear understanding of a host of additional issues. Most of these issues are tied up with the notion of what is meant by the term 'cognitive change' (for example, what counts as a genuine shift in our cognitive profile, relative to our current modes of cognitive processing?). It is also important to understand what is meant by the notion of the 'Web'. Many of the critical commentaries that have highlighted the putative hazards or shortcomings of the Web (Carr, 2010; Dreyfus, 2009) have tended to adopt a particular view of what the Web is. This view is dominated by our conventional modes of interaction with the Web – the use of browser-based forms of information access and the use of document- or page-centric modes of information representation and display. But why assume that our currently popular modes of Web-based interaction, and the technologies that support it, necessarily exhaust the range of possibilities when it comes to productive forms of cognitive engagement with the online world? The Web is something of a protean beast when it comes to user interaction and information access capabilities. New forms of information representation, new forms of user interaction technology, and new forms of application development all contribute to an ever-changing landscape against which our notions of the cognitive impact of the Web are always likely to be somewhat ephemeral.

What is the way we think?

Before we can conclude anything about the cognitive effects of the Web, or the Internet more generally, we need to understand something about the actual nature of human cognitive processing. This is important because it is easy to assume that the Web causes us to depart from a form of cognitive processing that is, in fact, largely uncharacteristic of the way cognition actually occurs in real-world situations. We might thus be tempted to assume that the standard form of human cognition is one of environmental detachment and inner contemplation – one in which the human thinker maintains a sustained focus on an internally-generated sequence of reason-respecting thoughts until some kind of cognitive outcome is achieved. This is the image that perhaps best captures our notions of what it means to think deeply about something, and it is the kind of image that is perhaps most vulnerable to the Web's influence. Rather than promoting our capacity for quiet contemplation and inner reflection, the Web seems to demand constant engagement with external sources of information. And rather than our thoughts being driven by association and reason, the Web encourages us to think in much the same as way we surf it: in a series of saccadic leaps between semantically-related (and sometimes not so semantically-related) resources. As we browse the Web, the sequential structure of our thinking becomes driven more by the lure of the links we see and less by the character of our thoughts themselves – our thoughts tend to follow the content of the links we click! Perhaps it is this fear of what the Web can do to our ability to think – our ability to engage in prolonged periods of environmentally-detached inner contemplation – that most fuels our fears about the cognitive implications of Web access. But the fear is justified only to

[8] See http://www.edge.org/q2010/q10_index.html.

the extent that the 'Rodinesque' image of human cognition – our image of human thinking as consisting in episodes of environmentally-decoupled reason and linear thought – is basically correct. If an alternative, more active and environmentally-engaged, model of human cognition could be advanced, the influence of the Web might not seem so malign.

The alternative image comes in the form of situated, distributed and embodied approaches to human cognition (Clark, 1999, 2008; Haugeland, 1998; Hutchins, 1995; Kirsh, 2009; Menary, 2010; Pfeifer & Bongard, 2007; Robbins & Ayded, 2009). These approaches have encouraged us to think about the way in which much of our cognitive success may be grounded in processing loops that factor in the contributions of our extra-neural social and technological environments. Consider the case of multiplying two three-digit numbers. One account of how we are able to multiply the two numbers might emphasize how we first derive some symbolic encoding of the visual (or auditory) input corresponding to the two numbers. It would then invoke a computational account according to which the inner symbols are manipulated in some way so as to achieve the correct mathematical outcome. Now contrast this with what is surely a more accurate (and ecologically-realistic) picture of how we implement long multiplication in the real-world. This alternative picture involves the active manipulation of external symbols in such a way that the kind of problem confronting the biological brain is profoundly simplified. In place of purely inner, environmentally-decoupled, computational operations we see a pattern of real world-involving perception-action cycles – ones in which single digit numbers are compared and intermediate computational results are stored in an external medium using (e.g.) pen and paper. This example, described in Wilson and Clark (2009), is a case of what we might call environmentally-extended cognition or cognitive extension (see also Clark, 2008). It takes what is, ostensibly, an inner cognitive capability (an ability to do long multiplication) and shows how crucial aspects of the problem-solving process can be (and usually are) delegated to aspects of the external world.

Another example of the way in which cognitive success may be predicated on our ability to actively engage with aspects of the external environment is the process of writing an academic paper or report. One view as to how we generate such artefacts might emphasize the role of purely inner resources in contributing to fully-formed thoughts, which are then serialized as words on paper. But this is seldom, if ever, how real academic texts get written. For better or worse, what generally tends to happen is that we start by writing down a few fragmentary thoughts and ideas, and these then prompt further thoughts and ideas. As the paper emerges, a variety of external resources, such as text and papers, often themselves heavily annotated with notes and marginalia, are continually consulted. As Clark (1997) argues:

> *"[the text] does not spring fully formed from inner cogitations. Instead, it is the product of a sustained and iterated sequence of interactions between my brain and a variety of external props. In these cases, I am willing to say, a good deal of actual thinking involves loops and circuits that run outside the head and through the local environment. Extended intellectual arguments and theses are almost always the products of brains acting in concert with multiple external resources. These resources enable us to pursue manipulations and juxtapositions of ideas and data that would quickly baffle the un-augmented brain." (pg. 207)*

What situated, distributed and embodied approaches to cognition provide, then, is an alternative image of human cognition. In place of the conventional image of the environmentally-detached human thinker, the alternative image emphasizes the role played by material embodiment and environmental embedding. Instead of cognitive success always being the product of inner reflection and quiet contemplation, we are afforded an image in which active engagement and interaction with the external environment plays an important role in at least some aspects of our daily cognitive activity. Obviously, there is no question here about our ability to actually engage in episodes of environmentally-decoupled reflection. We can, and clearly do, resort to this way of thinking on at least some occasions. The real issue is the extent to which this mode of thinking typifies much of our everyday cognition, and the extent to which we always need to engage in this kind of thinking in order to ensure cognitive success. If, as is suggested by theorists such as Clark (2008), our cognitive achievements are often reliant on forms of biotechnological and socio-cultural scaffolding, then we need to take this alternative view into account whenever we formulate responses to the question of whether the Web is changing the way we think. This is because the disruptive potential of the Web may seem greater (and far more injurious) to some forms of human thinking compared to others. If, for example, you subscribe to the view that human thinking is typified by quiet contemplation and sustained attention to a sequence of internally-generated thoughts, then the transition to a state-of-affairs in which human problem solving is dominated by extended bouts of active engagement with a set of external resources will seem like a turn for the worse. If, however, you view human cognition as something that characteristically involves the active recruitment of external resources into complex webs of extra-neural processing, then you will see the Web as simply providing the basis for additional forms of environmentally-situated cognitive processing. Importantly, our preconceptions about the nature of human cognitive processing – whether they have an internalist or externalist focus – are likely to guide our intuitions about the kind of cognitive change we see the Web producing and whether we see such changes as being, in general, a boon or a burden. If an ability for inner reflection is seen as the critical feature of our cognitive profile, and the mainstay of our intellectual and cognitive accomplishments, then the Web will be seen as, at best, influencing the kind of things we think about; i.e. the contents of our cognition. This may or may not yield cognitive advantages. However, in general, advocates of this view will be highly circumspect about the cognitive impact of the Web. They will tend to see extended periods of interaction with the Web as detracting from our opportunity to engage in 'real' thought, and they will regard any apparent deterioration in our ability to engage in environmentally-detached inner reflection as a sign of cognitive pathology. In particular, they are likely to regard anything that promotes a 'premature' curtailment of inner-directed thought as something that is potentially damaging to our cognitive capabilities. Instead of seeing the delicate inter-leaving of both inner- and outer-directed forms of attentional focus as a legitimate expression of at least some forms of continuous cognitive activity, they will tend to see any form of attention to external stimuli as reflecting a breakdown in cognition – as reflecting the harmful and unwanted intrusion of the external world into the inner sanctum of human thought and reason. In contrast to this view, advocates of more situated approaches to cognition are less likely to see our active engagement with the Web as merely providing the stuff of thought – the contents of cognition. For them, the Web may also play a role in shaping the process of thought – the mechanisms of cognition. This is an important shift in perspective, and it influences what we might expect to see in terms of the nature and scope of Web-based cognitive change.

The main point of this discussion is to show that the responses we might have to the question of whether the Web is changing the way we think is, in part, dictated by the kind of assumptions we make about the nature of human cognition. If we want to understand the cognitive effects of the Web, then we need to understand something about the actual nature of human cognition. We also need to be explicit about the kind of assumptions we are making regarding the potential for environmentally-extended forms of cognitive processing. To an "externalist" – one who adopts a situated, distributed or extended mind perspective – the cognitive effects of the Web will depend on the details of the informational contact the user is able to establish with Web-accessible resources. To an "internalist", the cognitive effects of the Web will be limited to whatever changes are seen in our capacity for environmentally-detached inner reflection and our ability to follow linear trains of thought.

When does a change in cognitive performance matter?

Whether or not we agree with the statement that the Web is changing the way we think, this says very little about the actual impact of the Web on our cognitive performance. A change in the way we think could be for the better or for the worse. Perhaps this is the most pressing issue in terms of the cognitive effects of the Web. The Web may be changing the way we think, but it is really the effect of the Web on our cognitive performance that is at the heart of our worries, hopes and fears.

One would think that it would be relatively easy to measure the effect of the Web in terms of our cognitive capabilities. We could, for example, take a sample of experienced and inexperienced Web users and subject them to a variety of cognitive tests, such as the tests of cognitive control employed by Ophir et al (2009). There is certainly some interest in running tests of this kind, although it is not entirely clear what such results by themselves will tell us about the causal role of the Web in effecting cognitive change. Another problem stems from the discussion in the previous section. If we accept that our contact with the Web may occasionally give rise to extended cognitive systems – systems in which cognitive processing is distributed across brain, body and world – then it becomes far more difficult to gauge the actual cognitive costs and benefits of Web access. The main problem here is that the cognitive profile of the socially- and technologically-integrated agent need not be that of the agent who is decoupled from her socio-technical surrounds. When an extended cognitive system emerges, the human agent forms an integrated system with some aspect of the external technological or informational environment, and the cognitive properties of that system are not necessarily those of the environmentally-decoupled agent. In most cases, the kind of cognitive achievements that are possible when the human agent is actively coupled with a technology are not necessarily possible when that agent is tested in isolation. To exemplify this, think of the cases of long multiplication and academic paper writing mentioned previously. If we remove the technological scaffolds of pen and paper, the human agent may find herself struggling to demonstrate the same kind of mathematical and literary competence as was possible when such resources were readily available. Now suppose that over the course of extensive use with pen and paper, the performance profile of the basic biological agent deteriorates. Perhaps the agent becomes so accustomed to using the external resources that she is no longer able to rely on the capacities of her bare biological brain. Should we say in these cases that the technological scaffolds have exerted a negative influence, that they have undermined the cognitive capabilities of the human subject, and that the use of such resources should be limited on the grounds that they are harmful to our cognitive well-being? This conclusion only seems to be warranted if we can convince ourselves that *our* cognitive performance profiles should always and only be judged relative to our biologically-bounded, brain-based achievements. To some of us, however, this conclusion does not seem entirely plausible. It may be that prolonged use of a technology produces changes in brain-based bio-cognitive functioning, but to say that the technology is thus exerting a negative impact on cognition depends very much on whether one accepts that the performance of the technology-wielding agent should *not* count as part of that agent's cognitive profile. To an extended mind theorist – one who accepts the possibility of environmentally-extended forms of cognitive processing – the performance of a properly integrated biotechnological system (comprising, perhaps, human agent and Web technologies) is a legitimate part of the cognitive profile of an agent. To such a theorist, the decision about the positive or negative contributions of a technology to our overall cognitive health will depend less on the details of bio-cognitive processing and more on the details of the bio-technological coupling (for example, whether the technology is always available to influence thought and action in cognitively-productive ways). In this case, it seems perfectly possible for a technology to make a net positive contribution to cognitive performance, in spite of a possible degradation in brain-based bio-cognitive capabilities, providing that the following conditions are met:

1. the technology does actually improve performance when it is being used (in particular, it supports levels of performance that would be difficult or impossible to attain if the individual was to rely on the operation of the bare biological brain);
2. the technology is readily available (in particular, the technology is always at hand whenever it is needed);
3. the technology is highly reliable (i.e. it generally performs as it is supposed to, certainly to a level that promotes trust in whatever outcomes it delivers); and
4. the technology does not negatively affect our ability to engage in other forms of cognitive extension (if the cognitive changes introduced by one technology interfere with the use of other technologies, then this may be a cause for concern).

Providing these conditions are met, we should not, I think, be overly alarmed by technology-induced changes in our bio-cognitive capabilities. What matters when it comes to human cognition is that we are able to engineer our technological environments in ways that alleviate much of the processing burden on the biological brain. As Clark (1997) comments, "Our brains make the world smart so that we can be dumb in peace" (pg. 180).

How significant are changes in brain function?

At first blush, the finding that exposure to the Web can cause apparent changes in brain activity (see Small et al., 2009) seems to provide compelling evidence of the Web's potential to alter cognitive functioning. The brain is, after all, the material substrate on which all the various elements of mind and cognition seem to depend, and it seems reasonable to conclude, therefore, that changes in the activity of the brain will result in a change in at least some aspect of our minds. Before we become overly enamoured with the results of brain imaging studies, however, we need to be clear what is and what is not revealed by such studies. For a start, it is not clear that all changes in neural activity are necessarily relevant to human cognition, or at least to the dimensions of human cognition in which we are interested. Given what we know about the neuroscience of human cognition, it seems likely that changes in higher brain regions (i.e. forebrain and midbrain structures) are more likely to affect cognitive performance as compared to lower brain regions. However, even if we do detect changes in the activity of some cognitively-relevant neurological system, it is difficult to know whether this reflects a fundamental change in some aspect of our cognitive profile, or whether it simply reflects the adaptation of the neural system to some new technology-using context. If, as was mentioned above, the brain is but one element of an environmentally-extended nexus of cognitively-relevant resources, then it seems likely that we will encounter changes in neural processing whenever we situate the brain within a new technological context. This change need not necessarily indicate any deterioration in cognitive performance; rather, the change may simply reflect the accommodation of some new external resource into an extended cognitive circuit – a case, perhaps, of the brain adapting its activity to best suit the features of the new cognitive niche in which it finds itself. Of course, any situation that involves interaction with a new resource is likely to introduce at least some changes in neural processing simply because new technologies often require the acquisition of new skills. The brain of someone who has never previously encountered a mouse and keyboard may need to incorporate a number of adjustments before skilful interaction with the Web can take place[9]. Similarly, we may expect to observe changes in neural functioning as a person evolves from a novice chess player into a chess grand master. Whether such changes reliably indicate a change in cognitive functioning is a moot point; however, they do not, by themselves, indicate that the cognitive potential of the human agent is either enhanced or undermined as a result of using a new technology.

The problem with brain imaging studies, like that of Small et al (2009), is that they actually tell us very little about the way in which a new technology is affecting cognitive processing, particularly if we embrace the notions of distributed, situated and extended cognition outlined above. Suppose we do observe a change in some aspect of neural functioning. Does this change indicate a fundamental shift in the way cognitive processing is performed? This seems likely only if we accept that all aspects of our cognitive profile are neurally-realized. If this is not the case, perhaps because some aspects of our cognitive processing are realized by environmentally-extended circuits, then it is possible that the same basic cognitive process is being implemented the same as before, just using different mechanisms (i.e. the mechanistic realization of the process has changed, but the nature of the process itself remains unchanged). On the other hand, suppose we fail to observe a change in neural functioning. Does this indicate that cognition is unaffected by the introduction of a new technology? Again, this only seems to be the case if we accept that all that matters to mind and cognition is what happens in the brain. If we accept the possibility that the machinery of mind and cognition extend beyond the traditional borders of skin and skull, then it seems that cognitive functioning could be radically altered by the ways in which specific aspects of the external environment are recruited as part of an environmentally-extended cognitive process, independently of whatever happens in the brain. As Clark[10] comments:

> *"Suppose we convince ourselves, by whatever means, that as far as the basic mode of operation of the brain goes, Internet experience is not altering it [cognition] one whit. That supports a negative answer only if we assume that the routines that fix the 'nature of human thinking' must be thoroughly biological:*

[9] Of course, it is difficult to know to what extent a lot of the research findings relating to Web-based changes in neurological or cognitive functioning are attributable to the kinds of user interaction technologies we rely on to access Web content. Conventional modes of Web access do not necessarily exhaust the range of potential interaction opportunities that future technological developments could bring (see below), and this has significant implications for how we think about the potential or actual cognitive effects of the Web.

[10] See http://www.edge.org/q2010/q10_5.html#clark.

> *that they must be routines running within, and only within, the individual human brain. But surely it is this assumption that our experiences with the Internet (and with other 'intelligence amplifiers' before it) most clearly calls into question. Perhaps the Internet is changing the 'way we think' by changing the circuits that get to implement some aspects of human thinking, providing some hybrid (biological and non-biological) circuitry for thought itself."*

The importance of brain imaging studies, such as those of Small et al (2009), are therefore difficult to assess in terms of their relevance to Web-induced changes in cognitive functioning. Such studies only really become of interest once we are able to detect a change in cognitive performance and are then able to relate the changes to changes in neural processing. Even in this case, however, it is often unclear whether the observed changes constitute a cause for concern for the reasons outlined in the previous section (see 'When does a change in cognitive performance matter?').

What do we mean by the 'Web'?

Suppose we accept the view that the Web *is* influencing cognition, and that, in general, the nature of the influence is negative. Perhaps we can assume that the worries and concerns expressed by Carr are indeed correct and that the Web is exerting a largely malign influence on our ability to think, read and remember. Suppose, for the sake of argument, that we also discover that the effects of the Web on our offline, bio-cognitive capabilities (i.e. the capabilities supported by the bare biological brain) are not compensated for by an overall improvement in our online, bio-technologically hybrid performance. Does this putative state-of-affairs mean that the Web is generally bad for our overall cognitive well-being?

In answering this question, I think it helps to pay attention to yet another question: what exactly do we mean by the 'Web'? The fact is that many people's concept of the Web is governed by their experiences with a particular set of technologies, namely HTML Web pages and Web browser technologies. There is nothing wrong with this view. Most of our interaction with the Web is based on the use of precisely these technologies. It is mistake, however, to assume that these kinds of technologies, and the interactive capabilities they support, exhaust the potential opportunities for the Web to exert some sort of cognitive influence (either positive or negative). This is why we should be cautious of blanket statements that either praise or pillory the Web on cognitive grounds. The current set of technologies we associate with the Web may yield a particular set of cognitive vices and virtues, but these do not necessarily speak to the cognitive impacts of future Web-based technologies. Later in the paper, we will encounter discussions of linked data (Bizer et al., 2009) and augmented reality, both of which provide alternative visions of what the Web is and the kind of capabilities it might support. The cognitive impact of such technologies may be either positive or negative; however, what is important for present purposes is simply that we should not be misled into thinking that a particular view of the Web (e.g. as browser-based access to HTML resources) necessarily exhausts the space of possibilities concerning future forms of Web-based cognitive change.

Extending the Mind: Cognitive Extension and the Future Web

The Web is a technology that is transforming many aspects of our lives, and if we are to believe the views of people like Carr, this transformative influence may very well extend to aspects of the human mind. For all that, however, I believe the cognitive effects of the Web, at least in its current form, are probably quite limited. True, some aspects of our cognitive profile may be enhanced or undermined as a result of Web access, but, in general, I suspect the Web is not causing any major shift in our cognitive power and potential. At least there is little hard evidence available at the present time to warrant such a conclusion.

This is not to say that the Web does not have the *potential* to effect major cognitive change. The Web is, I think, a technology that opens up all kind of opportunities for cognitive transformation. The realization of these transformative opportunities requires, I suspect, a significant change in the way we access online information. In particular, I think we need to move away from our traditional notions of Web-based information access and embrace an alternative vision of the Web. This alternative vision does not necessarily dispense with our familiar ways of interacting with the Web – the use of Web browsers, HTML documents and so on – but it does introduce a range of new interactive opportunities and possibilities, all of which are supported by a host of technological innovations.

My aim in this section is to sketch out this alternative vision of the Web and show why it is better placed to influence human cognition. The technological innovations that feature as part of this vision are those that I think are most likely to enable future forms of Web-based cognitive extension; i.e. the creation of extended cognitive systems comprising human agents, Web-accessible information resources and Web-enabled technological devices. In general, however, these technological innovations do not require any radical shifts in our current engineering capabilities, and, in many cases, early forms of the technologies are already beginning to appear. The vision of the future Web that emerges on the back of these technologies is a Web that is more suitable for participation in cognitively-potent biotechnological mergers. The new technologies will affect our access to online information in precisely those ways that are required for the emergence of extended cognitive systems, and the result will be a Web that is fundamentally more likely to transform our cognitive and epistemic profiles.

Writing and Thinking

In order to understand something about the transformative potential of a technology, it helps to focus on a specific example. Writing is a good example here because it involves the use of specific props, aids and artefacts to supplement our biologically-based modes of information storage, manipulation and retrieval. The invention of writing also arguably produced profound shifts in our cognitive power and potential. In particular, the invention of writing caused a significant change in the way we structured our intellectual activities, both individually and collectively. Writing enabled us to serialize our thoughts to an external medium and then revisit those thoughts at different times in different contexts[11]. Such modes of external storage also facilitated the spread of thoughts and ideas. They enabled the ideas of one person to influence the thoughts of many others. Of all the

[11] An important aspect of writing is that it enables us to revisit our thoughts at different points in time. Thus, we can approach our thoughts and ideas while in different mental (both cognitive and emotional) states and frames of mind. This is important because the kinds of experiences we have between the initial serialization of a thought and the subsequent reinstatement of that thought might enable us to form associations and insights that were simply impossible when we first had the thought. By revisiting our own thoughts (and indeed those of others) in different contexts, we can press maximum cognitive potential from those thoughts. Thoughts that failed to inspire productive shifts in our thinking at one point in time, can suddenly, by virtue of the new context, deliver new insights and open up new opportunities for cognitive and intellectual progress.

technological innovations that have had an influence on the way we think, none has perhaps been more profound than the development of writing.

As with any transformative technology, writing has had its fair share of critics. In Plato's Phaedrus, for example, Socrates expresses his concerns about the invention of writing:

> *"If men learn this, it will implant forgetfulness in their souls; they will cease to exercise memory because they rely on that which is written, calling things to remembrance no longer from within themselves, but by means of external marks. What you have discovered is a recipe not for memory, but for reminder. And it is no true wisdom that you offer your disciples, but only its semblance, for by telling them of many things without teaching them you will make them seem to know much, while for the most part they know nothing, and as men filled, not with wisdom, but with the conceit of wisdom, they will be a burden to their fellows." (Hackworth, 1952; pg. 157)*

Socrates' concerns about the use of writing to unburden our memories may, to some extent, have been well-founded. We clearly do rely on the external storage of information to support and (perhaps supplant) our bio-memory capabilities. However, concerns about the broader implications of writing on our cognitive capabilities do not seem particularly valid. Our intellectual heritage owes much to the invention of writing, and few, I suspect, would now advocate the abandonment of writing as a means to enhance our cognitive functioning. Even if it could be shown that writing has undermined our bio-memory capabilities, it seems unlikely that we would regard ourselves as better off, in a cognitive sense, if we were suddenly obliged to abandon writing and return to biologically-based modes of information storage and recall.

There are two reasons, I think, why Socrates' concerns about writing seem misguided, and both of these reasons touch on issues relevant to cognitive extension. The first reason relates to the easy availability, reliability and usability of writing technology. These factors, coupled with our ability to use writing technology effectively (an ability acquired though many years of formal instruction and education), means that we now regard *our* cognitive profile as simply a composite of whatever abilities are bestowed on us by both our biological brains and our ability to create, manipulate and exploit the written word. Whenever technologies become as readily available, as reliable, and as easy to use as most writing technologies, they promote forms of cognitive extension in which the cognitive abilities of the larger system (in the case of writing, the system comprising human agent, writing utensils and durable storage media) are simply seen as the cognitive abilities of the human agent. In such situations, we become so bonded with our technologies that they become a part of who we are as cognitive agents. The cognitive abilities of the brain, working in concert with the external technology, are seen as just as much a part of the cognitive agent as the cognitive abilities produced by the inter-operation of (e.g.) the hippocampus and hypothalamus. Just as it would not make much sense to say that *our* cognitive abilities are those of the hippocampus or frontal lobes minus whatever contributions are made by the hypothalamus, it does not make much sense to say (in some situations at least) that our cognitive profile is based solely on the operation of the biological brain minus a specific technology. In cases where a technology is readily available, and its use has become almost automatic for the human user (in the words of Clark (2007), it has become 'transparent in use'), then we are increasingly likely to see our cognitive capabilities as simply those engendered by the larger environmentally-extended cognitive systems in which much of our real world cognitive processing takes place.

A second reason why I think Socrates' views about writing fail to strike a chord with us is because his concerns centre on one specific aspect of writing, namely the use of writing technology to support the external storage of information. This is undoubtedly an important aspect of our writing capabilities, but it ignores the other kinds of roles that writing may play in supporting our thought processes. In particular, the storage-based view of writing fails to appreciate the active role played by the outputs of our writing in the driving of our thoughts. If we think of writing as merely the external storage of facts that can be reloaded into consciousness whenever the need takes us, then perhaps not much is gained in a cognitive sense. We simply alleviate our brains of some of their memory burden by offloading storage functions onto the external world. In general, however, the products of our writing activity fulfil more than this simple offloading function. They do not simply reinstate thoughts we once had. They also structure, constrain and influence our thinking processes. Recall the case of academic paper writing discussed earlier in this paper. What we see in the case of academic paper writing is a temporally fine-tuned interaction between the biological brain and a set of external resources, both of which influence one another in the form of a closed-loop feedback system. The brain generates some initial ideas which, when serialized to some bio-external media, serve as the trigger for additional (neurally-realized) thoughts and ideas. We don't just use the process of writing as a means of recording our thoughts and ideas. The writing process itself contributes to the kind of thoughts and ideas we have. In general, when we think about the cognitive impacts of writing (or of any cognitive technology), we should not be limited to thinking about the role played by writing in the long-term storage of the products of our thinking (what we might call its long-term memory function). We should also think about the active role of writing in driving the thinking process itself: in supporting the production of thoughts as well as their long-term storage[12].

The potential role played by writing in the forming of our thoughts has been appreciated by a number of theorists. In the introduction to his book entitled 'Supersizing the Mind: Embodiment, Action and Cognitive Extension', Clark (2008) recounts an exchange between the Nobel Prize-winning physicist Richard Feynman and the historian Charles Weiner in which Feynman argues for the importance of his notes in contributing to his thoughts. Rather than representing a mere record of his internal cogitation, Feynman seems to be arguing that the process of creating the notes and sketches is part of the cognitive work itself. Such notions are highly compatible with the notion of cognitive extension, and Clark (2008) goes on to suggest that the cycle of information flow and influence established by Feynman and his notepad plays a crucial role in the realization of Feynman's thinking:

> *"...I would like to go further and suggest that Feynman was actually thinking on the paper. The loop through pen and paper is part of the physical machinery*

[12] There is an issue here concerning the impact that different forms of writing technology have on our cognition. We can thus think about the differences in writing ability enabled by the skilled use of a keyboard compared to conventional pen and paper. The temporal profile of each is very different. Keyboards enable us to serialize our thoughts much more quickly (in some cases) than conventional writing utensils. Whether this difference in temporal profile actually helps or hinders us in terms of our literary accomplishments is something worthy of further consideration.

> *responsible for the shape and flow of thoughts and ideas that we take, nonetheless, to be distinctively those of Richard Feynman." (pg. xxv)*

What Clark is basically suggesting here is that writing is constitutive of thinking: that writing plays an active role in the realization of our thoughts, and that the machinery of cognition (at least in this particular case) extends to include not just the biological brain, but also the elements of the bio-external environment that make the creation of symbolic artefacts (written words) possible[13]. Thinking, from this point of view, is not something that occurs solely within the head. It is also something that can be spread across a variety of extra-neural and extra-corporeal resources. Thinking, as with other types of cognitive processing, is sometimes literally extended into the world outside the head.

What writing gives us, then, is a range of opportunities for cognitive extension that fundamentally transform the kinds of things of which we are capable, both at an individual and collective level. We do not need to fear the cognitive implications of writing technology; for even if it could be shown that some aspects of our bio-cognitive performance profile are negatively affected by our ability to write, it would still be the case that writing earns its cognitive keep in terms of the net benefits it brings us (both individually and collectively). What makes writing a force for good is not just that it facilitates our thinking during those times in which we are actually coupled to our writing-related artefacts (although this is undoubtedly important); it is also that the technologies we use for writing are so readily and reliably available. In the case of writing, we do not distinguish between the cognitive capabilities of the technologically-isolated human agent and the cognitive capabilities of the larger biotechnological system (comprising human agent and writing technologies). Instead, we simply see the cognitive capabilities of the human agent as those of the larger system. The cognitive profile of the larger biotechnological system is essentially incorporated into our cognitive schema – our sense of who and what we are as cognitive agents. And what makes this kind of cognitive incorporation plausible, on subjective, social and (possibly) scientific grounds, is the way in which writing technologies are poised to help us in our everyday cognitive affairs – whenever we need their help, they are always there, ready to assist us.

The Making of an Extended Mind

These days we can feel secure in the cognitive rewards writing gives us. But it wasn't always so – particularly in Socrates' time. The fact is that almost all forms of technologically-based cognitive extension require a prolonged period of technological, sociological and even neurological adaptation. In the case of writing, the features that make writing technologies so apt to participate in extended cognitive systems have come at the end of a long sequence of technological innovation and social change. The widespread availability of pen and paper, for example, did not happen overnight. It took many years before these resources were available in sufficient quantity for them to become a standard part of our everyday lives – part of the persisting backdrop against which our cognitive schema emerges. And although technological innovation and adoption are clearly a major part of the story, they are not the only things that need to be considered. We are not born skilled writers, capable of wielding a pen or working a keyboard. Rather, our writing abilities emerge over

[13] Clark is apparently not alone in highlighting the potential constitutive role played by writing. Something similar is, in fact, recognized by Carr (2010). In discussing Friedrich Nietzsche's use of a typewriter and the supposed effect this had on his writing style, Nietzsche is quoted as saying that "Our writing equipment takes part in the forming of our thoughts" (pg. 19).

the course of many years of instruction and training, often undertaken as part of a formal education. Pen and paper may be simple technologies, but their proper use and exploitation comes only at the end of a rather prolonged period of socially-scaffolded neurological adaptation and configuration. Biotechnological bonding does not necessarily come for free; we often need to teach our brains how to press maximal cognitive benefit from the technologies we use.

The history of writing also teaches us about the importance of social practices and conventions in enabling a technology to realize its full potential. In using writing technology, we follow a socially-accepted set of principles and guidelines governing the proper use of those technologies, and, over time, those principles have evolved to enable the technology to meet its designed purpose. The importance of these social conventions is apparent when we look at early forms of writing. These were not necessarily conducive to cognitive enhancement in the way that later forms were. In fact, early forms of writing were heavily influenced by the traditions and practices of the preceding era, within which information was communicated by purely oral means. Writing initially assumed the form of a continuous stream of text (known as scriptura continua), which was devoid of any of the conventional orthographic features (such as word spacing and punctuation) that we now accept as standard features of a written text. The reason for this particular form of writing seems to be based on the fact that writing was initially seen as a means to record and re-present orally communicated information. It simply never occurred to the early writers that the new system of recording thoughts and ideas could be used independently of the spoken word:

> *"It's hard for us to imagine today, but no spaces separated the words in early writing. In the books inked by scribes, words ran together without any break across every line on every page, in what's now referred to as scriptura continua. The lack of word separation reflected language's origins in speech. When we talk, we don't insert pauses between each word – long stretches of syllables flow unbroken from our lips. It would never have crossed the minds of the first writers to put blank spaces between words. They were simply transcribing speech, writing what their ears told them to write." (Carr, 2010; pg. 61)*

Scriptura continua therefore seems to reflect an intermediate stage in the transition from oral to written culture. It was a form of writing that was heavily influenced by previous forms of information communication, and like many innovations, it took time for it to become optimally suited to its target audience.

The making of an extended mind or extended cognitive system is not, therefore, necessarily straightforward. It sometimes takes a long time for a technology to be available in the right form, and we sometimes need to engage in extensive training before we can derive maximum cognitive benefit from its use. In addition to this, the use of a technology is often guided by social conventions, and these need to be carefully aligned with both the form of the technology and our ability to use it (or our ability to learn how to use it).

The Web: A Technology to Bond With?

Writing is a good example of a technology that has exerted a profound influence on our cognitive capabilities. Perhaps no other invention, other than the invention of language itself, can have exerted such an influence on our cognitive and epistemic potential both as individuals and as a society. And what is true of writing is also, I want to suggest, at least partly true of the Web. Like the invention of writing, the invention of the Web has provided us with a rich range of opportunities for cognitive extension and transformation, and, as is the case with writing, the transformative potential of the Web, both for society and ourselves, rests on a particular set of features concerning the availability, reliability and usability of the new technology. These features parallel those identified by Clark and Chalmers (1998) as part of their discussion of what makes an external information resource part an extended mind. In order for an information resource to legitimately count as part of an extended mind:

1. "…the resource must be available and typically invoked" (Clark, 2010). **[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, 2010). **[Trust Criterion]**
3. "…information contained in the resource should be easily accessible as and when required" (Clark, 2010). **[Accessibility Criterion]**

At present, it is unclear to what extent the Web, as currently constituted, possesses the kind of features that would support the emergence of environmentally-extended cognitive systems. The primary problem is that most forms of cognitive extension depend on a particular form of information flow and influence in which there is a close temporal coupling between the various elements that comprise the extended system. In the case of the extended mind, for example, it seems reasonable to insist that there should be some functional similarity in the influence exerted by both bio-external and bio-internal information sources. If external information should exert an influence on our thoughts and actions that is profoundly unlike that seen in the case of (for example) bio-memory, then it seems unlikely that we approach the kind of conditions under which genuine forms of mind-extending bio-technological integration take place. And if we think about the kind of informational contact we have with the conventional Web (the Web of HTML documents), then it seems unlikely that we will ever have a form of contact in which information can be accessed in a way that resembles the contents of our long-term memories. The nature of our contact with the Web seems to lack the kind of features that would make it the target of genuine forms of cognitive extension (see Smart et al., 2009).

We should not necessarily be surprised by this conclusion. The Web is a relatively recent technology, and it is still undergoing rapid development and change. As we saw in the case of writing, cognitively-potent forms of biotechnological merger do not come for free. They are sometimes the end product of a complex, often co-dependent, process of technological innovation, social change and neural configuration. The relatively recent emergence of the Web means that we should not expect it to immediately fulfil all of our requirements in terms of its capacity for cognitive extension. It sometimes takes time for the true transformative potential of a technology to be fully realized.

One specific way in which I think the history of writing informs our understanding of the current shortcomings of the Web (specifically concerning its capacity to support the emergence of extended cognitive systems) concerns the transition from early to modern forms of writing. We saw that early forms of writing, like the use of scriptura continua, were heavily influenced by the preceding oral tradition. It took some time before the social practices and conventions associated with writing were optimized to take account of its new cognitive role. Like early forms of writing, I want to suggest that the current form of the Web (the Web of HTML documents) is heavily influenced and informed by the practices and conventions of the era that preceded it: the era of the written and printed word. When information is published on the Web, it is typically done so in the form of 'pages' that communicate information in much the same form as we would expect to see in a printed document. This mode of information presentation is, like early forms of writing, not necessarily best suited to the emergence of extended cognitive systems. But perhaps the current version of the Web is, like scriptura continua, an initial form of a technology that is attempting to free itself of the metaphors of a previous era and evolve into something that is far more suited to the realization of its true cognitive potential. What the history of writing teaches us is that we should not mistake the early forms of a technology as constituting the final word in terms of that technology's ability to transform our cognitive capabilities.

Creating the Cognitive Machine

In general, much of the scientific interest concerning the effect of new Web-based technologies is concentrated at the social level. New technologies have been seen as enabling forms of social interaction and engagement that either reflect the dynamics of our existing social relationships or which give rise to new kinds of social processes. In one recent paper, Hendler and Berners-Lee (2010) discuss the implications of the Web in terms of the emergence of what they call 'social machines'. A social machine is a form of Web technology that works in concert with a human community to support interaction, collaboration and information sharing. According to Hendler and Berners-Lee, we are already witnessing the appearance of social machines in the form of social networking sites. However, the notion of the social machine is not just a way of conceptualizing the impact of a new technology on our social activities and processes. It is also a way of envisioning a particular kind of capability; in particular, one that demands the development and adoption of certain technologies. Hendler and Berners-Lee (2010) thus argue:

> *"This technology is needed because the magnitude of the problems that our society faces today are such that only the concerted effort of groups of people operating with a joint power much greater than that of the individual can hope to provide solutions. These problems, whether they be major issues such as curing cancer or addressing climate change, or more local like the creation of a community to address a social issue on the south side of Chicago, require people from a number of fields and backgrounds to work together. We believe that a new generation of Web technologies will be needed to address these issues..." (pg. 157)*

The notion of the social machine thus serves a dual purpose: it focuses our attention on the potential effects of the Web in terms of the creation of new forms of social interaction, collaboration and information exchange, and it also establishes a set of goals that can be used to guide current technology development efforts.

When it comes to considering the future technological evolution of the Web, it helps to think, not only about the social effects of the Web, but also about its cognitive effects. Just as new technological innovations may be seen as supporting the emergence of social machines, it is also useful to think about the way in which new technologies support the emergence of environmentally-extended cognitive machinery – how, in other words, they might give rise to what we could call 'cognitive machines'. As with the notion of social machines, the notion of cognitive machines serves a dual purpose. It acts as both a lens through which we can interpret some of the effects of the Web (specifically its effects on human cognition), and it identifies a set of capability targets that can be used to guide and orient technology development efforts. In particular, once we understand the kind of features that are required to support various forms of Web-based cognitive extension, we can use this understanding both to guide technology development and to anticipate the consequences of technology adoption.

The transition to a Web that is capable of supporting the emergence of extended cognitive systems – or cognitive machines – requires a number of innovations, and together these innovations give us an alternative vision of the Web and the kind of interactive opportunities it affords. To describe all aspects of this vision would take too long, but it is worth focusing on a couple of areas of technology development that are already beginning to change our relationship to Web-based information. The first of these areas concerns the use of linked data formats to change the way information content is represented on the Web (Bizer et al., 2009). Such formats improve both the accessibility (in terms of the retrieval of specific pieces of relevant information) and versatility (in terms of flexible modes of presentation) of information content. The second area of technology development concerns the use of new kinds of display devices and augmented reality capabilities. These change the nature of our relationship to information on the Web by making that information more accessible and more suitably poised to influence our everyday thought and action. Both of these areas of technology development and innovation are described in greater detail in the following sections.

The Missing Link: Towards the Web of Data

One relatively recent focus of scientific attention and technological innovation on the World Wide Web concerns the formats that are used to represent online information. Ever since the invention of the Web, the dominant way of representing information content has been via the use of HTML. Traditionally, information has been delivered in the form of HTML documents, which are accessed by browser technology and then presented to human users in the form of Web pages. This page, or document-centric mode of representation, has significant implications for how we access information content, and it affects the kinds of influence that Web-based information can exert over our thoughts and actions. If we want to enable the kinds of information flow and influence that support the emergence of extended cognitive systems, then we need to ensure that our contact with the Web fulfils the kind of criteria alluded to in the previous section (i.e. access, availability and trust). However, it is by no means clear that our current reliance on document-centric modes of information delivery do enable us to meet these criteria.

As an example of the shortcomings of document-centric modes of information representation, think about the problem of accessing factual information from a Web-accessible resource, such as Wikipedia. Even if the delays associated with document retrieval (i.e. 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, he or she cannot do this without reliably fixing the physical location of the information (perhaps by copying the required information to a local resource[14]). Ideally, what is required in order for Web-based content to count as part of an extended mind is that the relevant factual content should be available to guide thought and action in the way we have come to expect our thoughts and actions to be guided by information retrieved from bio-memory. The problem with the conventional, document-centric Web – the Web of Documents – is that the relevant pieces of information that are required to guide, scaffold and constrain our thinking are usually embedded in a mass of other distracting information. This makes it difficult to see how current forms of Web-based content could have the right kind of functional poise to count as part of our personal body of knowledge and beliefs (see Smart et al., 2009).

Fortunately, an alternative approach to the representation of online content is emerging alongside the conventional of Web of Documents. This is the Web of Data (Bizer et al., 2009), which is based on the idea that the Web should serve as a globally-distributed database in which data linkages are established by dereferenceable URIs. The transition from document-centric to data-centric modes of information representation is, I think, an important step in the technological evolution of the Web, particularly when it comes to the notion of Web-based forms of cognitive extension[15]. What is important for Web-based forms of cognitive extension are flexible modes of data integration, aggregation, filtering 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. The Web of Data supports these capabilities in a number of ways. In particular, it countenances representational formats that are:

1. largely independent of specific presentational formats or usage contexts (this supports flexibility in the way information content is presented; it also enables data to be used in different ways in different application contexts),
2. centred on the representation of limited sets of data or data items (this supports the rapid retrieval and presentation of specific pieces of information), and
3. semantically-enriched (this supports the retrieval of relevant information).

This mix of features brings us a step closer to establishing the kind of state-of-affairs associated with the emergence of cognitively-extended systems and the mechanistic realization of environmentally-extended mental states. The Web of Data is not necessarily the final step in this process, but it does provide an important link in the technological evolution of the Web, particularly as regards the future realization of cognitive machines.

The Real World Web

The gradual transition from document-centric to more data-centric modes of information representation is, I think, one of the ways in which the Web is evolving to provide new opportunities for cognitive augmentation and enhancement. However, without the correlative development of suitable interaction mechanisms, the potential for the new representational formats to genuinely transform our cognitive and epistemic potential is still somewhat limited. Although there are new

[14] 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.

[15] Paradoxically, although the aim of the linked data initiative is to provide content that is primarily suited for machine-based processing, while the aim of the conventional Web (the Web of Documents) is to provide content that is primarily suited for human consumption, it is the Web of Data that, I think, provides the better opportunity for human-centred cognitive transformation.

types of browsers that enable human users to browse the Web of Data, it is unlikely, I think, that these browsers will introduce any radically different forms of informational contact, at least relative to the kind of contact already afforded by conventional Web browsers. Rather than focus on the development of browser interfaces that simply take existing functionality and adapt it for the Web of Data, I suggest that we need to think about radically new forms of information display and user interaction. We need to move beyond the browser interface, which for too long has dominated our notions of informational contact with Web-accessible information resources. Instead, I suggest that we need to entertain a new vision of the Web; one which makes bio-external information resources suitably poised to participate in the emergence of environmentally-extended cognitive machines. Let us refer to this new vision of the Web as the 'Real World Web'.

As is suggested by its name, real world environments are at the heart of the concept of the Real World Web. The basic idea is that Web-based information should, wherever possible, be embedded in the real world and easily accessed as part of our everyday interactions with that world. Information about everyday objects should be associated with those objects, information about locations should be accessible in those locations, and information about people should be 'attached' to those people. In all cases, the information should be immediately accessible and easily processed. It must be able to guide thought and action in the way in which our everyday cognition is supported by the information retrieval operations of our own biological brain. What this means, in effect, is that information should be immediately accessible to our perceptual systems. It should require little or no effort to make the information available for perception, and, in most cases, the information should be delivered automatically, with little or no conscious effort required to make that information available.

This vision is one which modifies our traditional modes of interaction with the Web in a number of ways. Instead of the retrieval of relevant information being entirely the responsibility of the human agent, the notion of the Real World Web advocates a more intelligent and proactive Web: a Web which is capable of anticipating users' information requirements and making that information available in ways that support ongoing sequences of thought and action. It is also a vision that places the Web at the heart of our everyday embedded interactions with the world. Rather than information access requiring perceptual detachment and disengagement from our immediate surroundings (something that is required even with the most portable of mobile devices) the notion of the Real World Web seeks to make Web-based information access a standard feature of our everyday sensorimotor engagements with the world – it seeks to make the Web part of the perceptual backdrop against which our everyday thought and action takes shape. Finally, the notion of the Real World Web emphasizes a shift away from traditional browser-based modes of Web access, featuring the use of screen-based displays, keyboard-based interaction mechanisms and document-centred representational schemes. In place of conventional screen-based modes of information access, the Real World Web emphasizes the importance of more perceptually-direct forms of information access (e.g. the use of augmented reality devices to overlay Web-based information onto real-world objects and scenes). In place of conventional user interaction devices, such as mice and keyboards, the Real World Web advocates the use of alternative interaction mechanisms (more on which below); and in place of conventional document-centric modes of information representation, the Real World Web countenances a transition to more data-centric modes of information representation (see above). The main implication of this shift away from conventional browser-based modes of Web access is that we are enabled to see the Web in a new light: as something more than a passive set of information resources that need to be coaxed into useful cognitive service by our deliberate search and retrieval efforts. Our traditional modes of access with the Web encourage us to see the Web as something that is:

1. **passive:** we need to engage with the Web in a deliberate manner in order to retrieve relevant information. Information needs to be discovered, retrieved, filtered and interpreted; seldom does the Web support us in a proactive manner – providing the right information just when we need it.
2. **distinct from our everyday interaction with the world:** the Web may support our everyday decision-making and problem-solving behaviours, but, in general, this support comes at the cost of us having to divert our attention away from the problem at hand. Instead of information being immediately available to support our thoughts and actions, we are often required to suspend what we are doing in order to 'look things up' on the Web.
3. **impersonal:** the information content of the Web is, in general, not geared to specifically suit our idiosyncratic problem-solving needs and concerns; often we have to access information from several sources and adapt it for our own ends.

In place of this vision, the Real World Web gives us a vision of the Web as something that is proactive, personal and perceptually-immediate. Once we are afforded immediate perceptual access to Web-based information, and once such information becomes available at just the right time to support our goals, interests and concerns, then such information becomes, I suggest, far more capable of fulfilling the kind of conditions that merit the emergence of cognitive machines.

The Real World Web, it should be clear, relies on a rich range of sophisticated technical capabilities, most of which are, as yet, either unavailable or not in widespread use. This might be perceived as grounds for pessimism about the tenability of the Real World Web vision. However, for the most part, the kind of technological innovations required to make the Real World Web a reality are not beyond the reach of our current engineering capabilities, and, in some cases, early forms of the technologies are already starting to appear. One of the most interesting and relevant areas of recent technological innovation concerns the development of a range of highly portable augmented reality or mixed reality solutions. These are available in a variety of forms, from handheld mobile devices that overlay information onto a real-world scene via a camera and screen display, through to wearable-computing solutions, such as head-mounted displays and retinal projection systems. Most of these technologies are still the subject of ongoing research efforts; however, in most cases, the capability targets are clearly identified, and early prototypes of the systems are beginning to become available.

One research program that is seeking to develop portable projection systems that present information directly to an individual's field of view is the DARPA-funded ULTRA-Vis program[16]. This program builds on the capability vision provided by mobile device eyewear systems[17], which 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. Another focus of

[16] http://www.darpa.mil/ipto/Programs/uvis/uvis.asp
[17] http://www.microvision.com/wearable_displays/index.html

current research attention concerns the development of advanced retinal projection systems that project information directly onto the retina of the eye. One such system, which is being developed by researchers at the University of Washington, uses a set of microfabrication techniques to incorporate display micro-devices into a contact lens (Ho et al., 2008; Lingley & Parviz, 2008; Pandey et al., 2009). The contact lens is worn like any other contact lens, and it provides a see-through display that is both remotely powered and controlled via a wireless link. The ultimate promise of such devices is that they enable network-accessible information to be superimposed on real world objects and scenes, enhancing the kind of informational contact we have with our online world, and significantly enriching the range of perceptual cues, prompts and affordances that guide our everyday thoughts and actions.

Unfortunately, the capabilities provided by these augmented reality/mixed reality technologies are only part of what is required to realize the vision of the Real World Web. In general, the technical realization of the Real World Web rests on research and progress in the following areas:

1. **Augmented reality display devices.** Considerable research effort has already been devoted to these kinds of devices. They assume the form of head-mounted displays, virtual retinal displays and augmented reality applications for handheld mobile devices.
2. **Context-sensitive information retrieval and display capabilities.** A key capability for the Real World Web is the ability to proactively support the user in his or her everyday engagements with the world. This requires that information should be retrieved in a manner that is highly sensitive to the ongoing needs, interests and concerns of the human user. In particular, it means that information retrieval and presentation should be sensitive to various kinds of contextual information, such as the location of the user, the kind of cognitive activities in which they are engaged, and the kind of objects with which they are confronted. Context-sensitive modes of information retrieval and presentation are the focus of a number of ongoing research efforts (Bahrami et al., 2007; Rhodes, 2003; Rhodes & Maes, 2000), and further progress in this area will probably require significant advances in our ability to detect particular kinds of contextual information (for example, locations, objects[18], people[19] and situations[20]), as well as our ability to actually represent that information in a way that meliorates the information retrieval process (see Bao et al., 2010).
3. **Data-centric modes of information representation.** The use of linked data formats to support the realization of the Web of Data was discussed earlier in the paper. In general, the Web of Data will provide an information space that is much more suited to the Real World

[18] Improved object detection and recognition capabilities could come about as a result of advances in signal processing capability (including auditory as well as visual scene analysis capabilities), or they could stem from the use of interventions that simplify the detection/recognition problem (e.g. the use of radio-frequency identification tags).

[19] Person identification is being supported by progressive improvements in our ability to use a variety of kinds of biometric information. This includes the recent development of human gait recognition (e.g. Nixon & Carter, 2006) and ear morphology processing (Hurley et al., 2008) techniques, as well as the use of more conventional face processing technologies.

[20] In general, the ability to detect specific situations requires the integration of multiple types of contextual information, and it may therefore be somewhat more difficult to perform compared to the detection/recognition of particular objects. One example of the kind of research being undertaken in this space is the attempt to monitor and infer mission status information from both physical and contextual cues (Poltrock et al., 2008). This research is being undertaken as part of a larger research effort to provide more proactive forms of information assistance to military coalition teams.

Web than the current Web of Documents. In particular, the representational formats adopted by the Web of Data support the targeted retrieval of specific pieces of task-relevant information. They also provide the basis for highly flexible modes of information display, enabling information to be presented as simple textual prompts and alerts, as simple graphic elements, or as more complex graphic overlays on a visual object or scene.

4. **Flexible modes of information display.** Flexible modes of information display are important because we will often want to present information in a way that best supports real-world decision-making. Sometimes this may involve the presentation of information in text-based formats; at other times we will need to make use of graphic objects that guide thought and action in particular ways (for example, our interaction with physical objects could be guided by the use of virtual visual affordances that supplement those already associated with an object). Obviously, it will be important to avoid situations where the human user is overwhelmed with information, so the use of intelligent information filtering solutions (particularly those that are sensitive to the prevailing context) will be necessary in addition to advanced information display capabilities[21].
5. **Low cost[22]/minimal effort information retrieval solutions.** A key element of Web-based forms of cognitive extension is likely to be the advent of low cost/minimal effort solutions to the problem of interacting with online resources. What this means, in effect, is that it should be possible to implement information retrieval operations in a way that is very rapid and which requires the minimum of physical effort. One promising line of research here involves the development of sub-vocalization techniques for Web-based information retrieval (Jorgensen & Binsted, 2005). This allows users to silently voice instructions and commands that could be used to control a variety of forms of interaction with the Web; for example, it could enable users to pose questions in natural language which could then be used to direct information retrieval operations[23]. Another important strand of research and development concerns what are commonly referred to as neuro-interactive or brain interface technologies (Mason et al., 2007; Pfurtscheller et al., 2007)[24]. These devices are based on the idea that by processing the electrical activity of the brain, we can establish a direct link between neural activity and environmental control, enabling our thoughts (or at least the neural signals associated with our thoughts) to control at least some aspects of our physical engagement with the world (see Nicolelis, 2001).

As was mentioned previously, many of these capabilities are already the focus of ongoing research and development efforts, and there is, as such, a significant overlap between the goals of many

[21] One interesting line of research here concerns the use of subliminal cuing strategies to minimize the cognitive load imposed on users (DeVaul et al., 2005).

[22] The notion of 'cost' here is intended to reflect the amount of time and physical effort that is required to access information; it does not refer to the financial cost of accessing information. Information access cost has emerged as an important construct in terms of the specific mix of resources (bio-memory, motor actions, shifts of attention) that are recruited to solve a problem (see Clark, 2008; pg. 119). In particular, it has been shown that low information access cost (measured in terms of the time taken to access information) promotes the adoption of a solution strategy that delegates much of the information retrieval burden to perceptuo-motor routines rather than the use of bio-memory resources (Gray & Fu, 2004).

[23] One associated capability here concerns the translation of natural language queries into more formal query languages, such as those used to retrieve information on the Web of Data. A number of techniques and technologies are emerging to support this translation capability (see Smart, 2008).

[24] For a commercial example of such devices see http://www.emotiv.com/.

extant technology development programmes and the kind of technological advances required to support the technical realization of the Real World Web.

The notion of the Real World Web serves as a potential antidote to at least some of the criticisms made by commentators such as Dreyfus (2009) and Carr (2010). Dreyfus (2009), for example, admonishes those who hype the benefits of the Web for such things as distance learning. Such claims, he argues, ignore the importance of embodied forms of interaction with the real world which serve to mediate many of our advanced learning experiences. Dreyfus is right, I think, to draw attention to the limits of a technology that distances itself from the facts of material embodiment and environmental embedding, but such criticisms should only serve to remind us of the limitations of the Web as it is currently constituted. Future forms of engagement with Web-based information, such as that envisioned by the notion of the Real World Web, need not be so remote from our everyday embodied interaction with the World, and, indeed, much of the power and potential of the future Web, from the perspective of cognitive enhancement, may stem from the fact that online information is much more productively poised to shape, scaffold and constrain our everyday embodied interactions the world.

Carr (2010) also criticizes the Web in terms of our traditional browser-based access to online content. “A screen-based reading behaviour is emerging,” Carr claims, “which is characterized by browsing and scanning, keyword spotting, one-time reading and non-linear reading.” In fact, many of Carr’s concerns about the cognitive impact of the Web seem to be related to the current way in which we access online information and interact with the Web. It is not clear, for example, to what extent the putative effects of the Web on our ability to think, read and remember are purely the result of our currently popular forms of engagement with the Web (i.e. our use of conventional user interaction devices, HTML-based modes of information presentation and screen-based modes of information display). The important point, for present purposes, is simply that we should not see the properties of the current Web as defining the limits of its cognitive transformative potential. As we move into an era of wearable computers, augmented reality display devices, and Web-enabled objects[25], so we encounter new opportunities for cognitively-potent forms of engagement and interaction with the largest space of knowledge and information our species has ever seen.

Creativity and Collective Cognition

Thus far the discussion has centred on how the cognitive profile of individual human agents might be affected by current or future forms of the Web. There is, however, another aspect to this discussion that we haven’t touched on as yet. This is the role played by the Web in supporting episodes of collective or collaborative cognition. Ever since the advent of Web 2.0, which is characterized by greater levels of user participation in the creation, maintenance and editing of online content, the Web has provided ample opportunities for socially-distributed forms of information processing. Whether we choose to regard these forms of processing as genuinely cognitive or not (see Smart et al., 2010c), it is clear that the Web has opened up new ways in which both the contents and character of our cognitive activity can be influenced by other human agents.

[25] The notion of Web-enabled objects (or the Web of Things) is simply the idea that everyday objects are integrated into the Web and are therefore Web-accessible. This means that objects may publish information about themselves and their current status on the Web, and that, at least in some cases, objects may be manipulated via Web-based forms of communication (see http://www.webofthings.com for more information).

One way in which we might derive cognitive benefit from the Web is in terms of our collective creative potential. The Web facilitates the distribution of thoughts and ideas and these can, on occasion, stimulate the production of new thoughts and ideas. Our collective capabilities might be enhanced in this situation by virtue of the fact that exposure to the same idea will not necessarily produce the same cognitive response in all individuals. Individuals differ in terms of their expertise and experience and this may give rise to different kinds of cognitive responses. The widespread distribution of information, in this case, may be expected to stimulate creativity – enabling individuals to respond in highly diverse ways to the same body of information and producing net gains in the total number of generated thoughts and ideas.

The potential of the Web to support creative insight and intellectual progress is well recognized in the scientific community. As Stevan Harnad (1999) rightly notes, the Web allows us to accomplish something akin to 'scholarly skywriting' – scientific theories, thoughts, ideas, experimental results, and sometimes data, are made available in ways that are increasingly accessible to fellow academics and scientific colleagues. It is almost as if the outputs of scientific and intellectual enquiry were written in the sky for all to see.

Such notions highlight two important aspects of Web-based forms of information distribution and idea propagation: global reach and immediate access. Just as if ideas had been written in the sky, the use of the Web as a distribution medium means that our ideas have the potential to reach a very large number of people. Furthermore, once our thoughts and ideas are published online, they are available for immediate access and can exert an influence on others with a speed that is unrivalled by conventional print media. Skywriting, suggests Harnad (2004), has a temporality that is akin to the speed of thought. The Internet, he suggests, supports rates of information flow and influence that parallel those seen in the oral tradition. It enables information to be exchanged at roughly the speed at which most of our thinking takes place: the speed of speech:

> *"There is every reason to believe that our talking heads and their interacting minds will be incomparably more fecund once those lazy iterative cycles by which our knowledge had been created and cumulated in the Gutenberg era are restored to the speed of stone-age thought by skywriting in the PostGutenberg Galaxy." (Harnad, 2004)*

These two features – widespread distribution and rapid dissemination – make the Web hugely beneficial to our collective creative potential. Widespread distribution enables one idea to reach the minds of many others, thereby leading to productivity gains in the form of a greater number of diverse ideas. Rapid dissemination improves the rate at which new ideas get generated, thereby leading to efficiency gains.

The problem with this optimistic image is that the mode of information distribution promoted by the Web does not always seem of great benefit to our collective cognitive endeavours, especially when it comes to the generation or discovery of novel solutions (Huynh et al., 2010; Lazer & Friedman, 2007; Mason et al., 2005; Smart et al., 2010b). In our own research, for example, we have discovered that rapid and widespread modes of information distribution are not always the most effective for all types of collective problem-solving. In some cases, these modes of information distribution can actually stultify rather than stimulate collective creativity (Huynh et al., 2010; Smart et al., 2010b). What appears to underlie this deficit is the fact that rapid and widespread modes of information distribution tend to encourage lots of agents to prematurely settle on a particular solution, and this effectively prohibits the pursuit of alternative solution paths that may ultimately lead to better long-term outcomes.

Findings from the social psychological literature seem to complement these findings. Thus, one well known phenomenon in group problem-solving situations, specifically brainstorming sessions, is production blocking (Diehl & Stroebe, 1987). Production blocking is the tendency for the contributions of one individual to block or inhibit contributions from other group members, resulting in an reduced level of creativity relative to what might have been expected if the group members had worked independently. It thus seems that, at least in some situations, the tendency to share information can undermine the collective creative potential of a group of agents. Instead of stimulating a greater number and diversity of ideas, precipitant forms of information sharing can sometimes impede rather than improve the creative process.

What we arrive at, therefore, is a dichotomy in our views concerning the potential impact of the Web on collective creativity. On the one hand, it seems that the information distribution opportunities afforded by the Web stand to improve collective creativity. The Web enables us to distribute ideas to a global audience, and because we have immediate access to ideas as soon as they become available, it seems that the Web is well placed to deliver both productivity and efficiency gains when it comes to the generation of new ideas. On the other hand, it seems that the Web could potentially undermine our collective creative capabilities. The Web supports modes of information distribution that are both rapid and widespread, and these may lead to reductions in the number of unique ideas that get generated.

At this point, critics will be keen to point out that although the Web does potentially support rapid and widespread modes of information distribution, these are seldom, if ever, achieved in reality. The sheer size of the Web, it will be argued, means that any new idea will have limited distribution simply because it will tend to get lost in the noise of other competing information. Being able to publish your ideas in the sky does not necessarily increase their distribution if everyone else is publishing to the same place. Each idea will at best have limited uptake by the members of a community, and the chances are that it will processed at very different times. The end result will be that no one working on a particular problem is going to be exposed to exactly the same information at exactly the same time. There will always be some differential exposure to particular subsets of information.

This is an interesting line of defence because the size of the Web has often been seen as a source of problems relating to information overload and information retrieval. It probably never occurred to anyone to see the size of the Web as a means of obviating modes of information distribution that might otherwise prove pernicious to collective cognitive processes. The problem, of course, is that the difficulties posed by the Web's size (in terms of the problems it presents for information access and distribution) have been the focus of considerable recent research attention and technological innovation. We now have search engines, recommender systems, tagging systems, Google alerts, RSS feeds, Twitter recommendations, taxonomies, collaborative filtering capabilities, folksonomies, ontologies, social networking sites[26] and so on, all of which aim, in one way or another, to put us in

[26] In all likelihood, the small world nature of our social networks works in concert with these technologies to support the rapid and widespread distribution of information.

touch with specific subsets of information. When these information filtering technologies become widely adopted by a particular community, they serve to support precisely the kind of information distribution profile that might prove problematic for collective creativity. They ensure that all members of a particular problem-solving community – members of a distributed scientific team, say – all get exposed to the same ideas at the same time. The specific worry about this mode of information distribution is that it may very well encourage a community to come to a premature consensus about some problem, situation or other state-of-affairs, and thereby limit the pursuit of more productive intellectual trajectories.

Perhaps the key point to note here is that we should not necessarily assume that a particular mode of information distribution is always beneficial for cognitive processing. Neither should we assume that technologies that supposedly help us locate relevant information, such as search engines, collaborative filters, and so on, are always beneficial (or even benign) in terms of their cognitive consequences. In some cases, these filtering technologies may not so much broaden our intellectual horizons (by exposing us to a global space of diverse thoughts, ideas and views) as restrict our attention to a particular part of the information space, namely, that part that has already been the focus of lots of other individuals' browsing efforts[27]. When it comes to the increasing availability of online scholarly information content, Evans (2008) notes that electronic modes of information distribution have not necessarily resulted in academics being exposed to more ideas. Instead Evans argues that the transition to online modes of distributing scientific content is accompanied by a progressive narrowing in the focus of the academic community. At the heart of this phenomenon is the fact that the online world often channels our search efforts to particular places on the Web – as we follow a particular set of links, we end up following the same well-worn intellectual paths and trajectories that have been traversed by many of our peers. As Carr (2010) notes:

> *"Automated information filtering tools, such as search engines, tend to serve as amplifiers of popularity, quickly establishing and then continually reinforcing a consensus about what information is important and what isn't." (pg. 217)*

Thus, although the vision of scholarly skywriting is one in which researchers are rapidly put in touch with lots of diverse ideas and scientific findings, it is possible that our attempts to harness the power and potential of the Web have inadvertently led to a narrowing of scientific focus and attention, reducing our ability to maximally exploit the creative opportunities that the Web provides.

To summarise, it is really too early to tell what the precise effects of the Web are on our collective creative potential. Much more fundamental research work needs to be done, and we are only just beginning to understand how cognitive processing is influenced by the dynamics of information flow and influence in distributed network environments. One notion that may be useful in our understanding of this relationship is the notion of adaptive coupling (see Smart et al., 2010c). Adaptive coupling is simply the idea that the various elements of a cognitive system need to establish time-variant patterns of functional connectivity with one another such that the overall information processing performance of the entire ensemble (the cognitive system) is optimized. In this case, we need to understand how systemic performance is affected by different temporal profiles of information flow and influence, and, in all likelihood, different kinds of cognitive

[27] Thus, for any given community, collaborative filtering systems and algorithms like Google's PageRank algorithm may function to limit attention to a relatively small number of highly popular resources.

processing will demand different kinds of temporally-specific functional coupling. In the case of creative insights, for example, it may be important to introduce inefficiencies into the information distribution process, or it may be important to ensure that not all individuals get exposed to the same information at the same time. Perhaps one way in which this can be accomplished is by diversifying the kinds of recommender systems, filtering capabilities or Web search algorithms that get used to retrieve information. Thus, rather than encouraging everyone to adopt exactly the same information retrieval solution, it may be important to encourage the use of different tools, algorithms and classification systems[28]. It may also be important to recognize the value of occasionally removing ourselves from outside sources of influence – effectively disconnecting ourselves from the network-mediated flow of information and ideas. One of the elements of creativity is incubation (Wallas, 1926) – the period during which a problem is put aside, thereby allowing non-conscious processes to productively operate on problem-relevant information. It is during these periods of disconnection that we perhaps enable the incubation process to get a thorough foothold, and, therefore, perhaps a judicious mix of engagement and disengagement with online resources – the essence of the adaptive coupling thesis – is precisely what is needed to maximally exploit the cognitive potential of the World Wide Web.

Cognition and the Web: An Agenda for Web Science

The Web shows all the signs of being a technology that, like the invention of writing and the mechanical clock, is poised to transform both society and ourselves. Given the Web's transformative potential, and the important role it already plays in so many of our daily activities, the scientific study of the Web, and its relationship with wider society, has recently been proposed as the basis of a new scientific discipline called Web Science (Berners-Lee et al., 2006; Hendler et al., 2008; Shadbolt & Berners-Lee, 2008). Web Science treats the Web as its main focus of scientific study, but its remit extends beyond a simple analysis of the technical underpinnings of the Web. Web Science recognizes the complex relationships that exist between technology, society and individuals. It recognizes that technological innovation on the Web may impact society, perhaps giving rise to new forms of social interaction and engagement. Web Science also recognizes that technological innovation takes place against a complex sociological and psychological backdrop: a backdrop that includes, among other things, moral and ethical codes, legal constraints, social conventions and cognitive capabilities. Web Science is not, thus, purely techno-centric in focus. It recognizes that we are shaped by technologies and that we in turn shape the technologies we use.

Given the importance attached to understanding the bidirectional nature of the influences that exist between the Web and its users, Web Science is, by its very nature, highly interdisciplinary in focus. This makes it ideally suited to examining the complex inter-relationships that exist between human cognition and the Web. Web Science is also a discipline that combines analytic and synthetic

[28] Typically, the focus in many classification-oriented modelling activities is to establish a common set of community-wide constructs that can be used for resource classification purposes. Thus, in ontology engineering, the aim is usually to create a single ontology that reflects consensual views as the conceptual structure of some domain of discourse. In cases where multiple, distinct ontologies arise, perhaps as a result of divergent views and perspectives (see Smart & Engelbrecht, 2008), a range of ontology reconciliation techniques have been developed to support the mapping and merger of the ontologies (Kalfoglou & Schorlemmer, 2003; Noy, 2004). The result of these reconciliation techniques is to establish a unitary view of a domain that may (at least in some situations) function to further limit the diversity of resources a particular community is exposed to.

approaches. It seeks to develop an understanding that is grounded in empirical analysis, but it also aims to use that understanding as the basis for influencing the Web's development. Web Science thus provides an ideal interdisciplinary context in which many of the issues raised in this paper can be pursued in greater detail. It enables us to study the psycho-cognitive impact of the Web in its various forms (the Web of Documents, the Web of Data, the Web of Things, and so on), and it enables us to develop future Web technologies that best support cognitive processing at both the individual and collective levels.

The following highlights some strands of Web Science research that could be undertaken in order to improve our understanding of the relationship between human cognition and the Web. This is by no means a complete list. Notable omissions include the design and implementation of various technologies (e.g. applications, devices, protocols, representational formats) that are required to maximise the cognitive benefits of the Web. Another omission relates to the philosophical issues concerning Web-based forms of cognitive extension and collective intelligence (Halpin et al., 2010; Smart et al., 2009).

1. **Cognitive Effects of the Web.** This strand of research aims to investigate the cognitive effects of the Web; in particular, the effect the Web has on aspects of our offline cognitive performance profile. One specific line of enquiry concerns the performance of both 'heavy' and 'light' Web users on a range of psychological tests, such as those used by Ophir et al (2009) to measure cognitive control. We saw in the case of Ophir et al (2009) that heavy media multitaskers exhibited a cognitive profile that was characterized by bottom-up attentional control strategies. It will be interesting to see whether the heavy use of the Web is associated with a similar cognitive profile.
2. **Subjective Effects of Reliable Information Access.** We saw that a key aspect of the vision of the Real World Web concerns our immediate and reliable access to information. It will be interesting to see what effect this kind of access to information has on our subjective sense of our epistemic capabilities. For example, if a particular body of information was always available to us whenever we needed it be so, would we start to experience a shift in our epistemic relation to that information? Would we start to feel as if that external body of information was just as much a part of our own personal body of knowledge and beliefs as the contents of our bio-memories? The key issue here concerns the factors that influence our sense of when we are in possession of a particular body of information. What factors conspire to make us feel as though we 'know' something[29]? Is it something to do with the way in which our semantic memories are implemented, or is it something to do with the

[29] There is a side issue here concerning whether our sense of what we know is informed by the same processes that determine the content and character of much of our conscious experience. According to one influential theory of consciousness, our subjective perceptual experiences stem from an implicit knowledge of what are called sensorimotor contingencies or sensorimotor dependencies (i.e. our knowledge and expectations concerning the effect of movement or change on patterns of sensory stimulation) (Noë, 2004; Noë, 2009; O'Regan & Noë, 2001). It is interesting to contemplate whether our sense of what we know may be governed by an implicit knowledge of what I call 'epistemic contingencies' (i.e. our knowledge and expectations about the way in which specific kinds of information are made available by particular thoughts and actions). In the same way that our subjective sense of the visual scene is informed by our knowledge of how visual information reliably changes when we (e.g.) move our heads, our sense of what we know may be informed by our knowledge of how (linguistic?) information reliably becomes available to us whenever we engage in particular thoughts and actions, or whenever we find ourselves in particular situations.

mere accessibility and availability of information – the fact that the things we know are (usually) poised to influence our thoughts and actions irrespective of where we are and what we are doing? If our sense of what we know is limited by the availability of information, then it seems that a transition to something like the Real World Web may promote a profound shift in our sense of what we feel we know. If all the information on the Web was made available to us in a way that was similar to the information made available by our bio-memories, would we begin feel as though we knew the entire contents of the Web? And, irrespective of any subjective changes, when does it make social and scientific sense to say that we genuinely know something? Could the boundaries of our semantic memories coincide with the information contents of the Web? If not, why not?

3. **Web Content and Cognitive Processing.** Building on the results of studies investigating the effect of hypertext on memory and comprehension processes, it will be important to understand how the various features of Web content relate to cognitive processing. One strand of research that could be undertaken here seeks to more fully understand the relationship between hyperlinks and aspects of cognitive performance. For example, are the disruptive effects of heavily hyperlinked resources mediated by an increase in cognitive load (see DeStefano & LeFevre, 2007)? A second issue for research concerns the relationship between Web-based multimedia content and cognitive function. For example, does the existence of multiple information streams on a particular Web page disrupt a user's comprehension or memory of the page's content[30]? Finally, as we move into an era where linked data content is becoming more widely available on the Web of Data, we need to better understand the cognitive effects of our access to linked data resources. In particular, we need to know how the presentational and interactive mechanisms developed for the Web of Data affect memory and comprehension processes.
4. **Collective Cognition and the Web.** The effect of the Web on collective cognitive processing is poorly understood. We need to understand whether the advent of the Web is a boon or burden for such things as collective creativity. Of particular interest is the question of whether our use of information filtering technologies, like search engines, limits the exposure of a particular community to a particular set of facts and whether this contributes to the kind of phenomena seen in group problem-solving situations; for example, production blocking (Diehl & Stroebe, 1987) and the common knowledge effect (Stasser & Titus, 1985).

Conclusion

Throughout history, technologies have exerted a significant influence on humanity. Some technologies have contributed to profound forms of socio-economic change, fundamentally transforming the nature of social activities and the structure of social relationships. Others have contributed to a profound change in our cognitive profiles, fundamentally transforming the way we think about the world, modifying the basic character of our cognitive processes, and occasionally extending the reach of our cognitive capabilities. Like the transformative technologies that preceded

[30] One reason to think that multiple information streams may negatively impact a user's memory of Web page content stems from a study investigating the effect of visually complex environments on factual recall. Bergen et al (2005) investigated whether the visually complex environment used by CNN during news broadcasts negatively affected college students' ability to recall news-related facts. They showed that factual recall was worse in the standard visually complex environment as when compared to a more simple environment consisting of the news anchor alone with all other information items (e.g. news crawls, stock quotes, weather forecast icons, etc.) removed.

it, the Web is a technology that is potentially poised to influence both society and ourselves. At present, much of the empirical research related to the Web has focused on its social and societal impact: the effect the Web is having on social, political and cultural processes. However, the cognitive impacts of the Web are also an increasing focus of scientific interest and research. Future empirical research should improve our understanding of the precise impact of the Web on our cognitive functioning, and this understanding will serve to guide the development of future Web-based technologies.

Whatever the outcome of future empirical studies, we need to bear in mind that the Web is a complex evolving system. Our views about its cognitive impacts are, in all likelihood, historically specific. Even if it could be shown that the current version of the Web is exerting a largely negative influence on our cognitive capabilities, this by no means implies that future versions of the Web will fail to deliver cognitive benefits. As we saw in the case of writing, some technologies take time to become ideally suited to support cognition, and perhaps something similar is true of the Web. As was the case with writing, our ability to gain maximal cognitive benefit from the Web will depend on a progressive optimization of the technologies we use and the social conventions that dictate how we use them. The current Web is surely an important point in our cultural and intellectual history, but it need not be the final word in terms of our cognitive and epistemic transformation. Perhaps, with time, we will come to see the current Web as laying the groundwork for more potent and profound forms of cognitive enhancement – an important, albeit temporary, milestone en route to our Web-enabled cognitive future.

References

Bahrami, A., Yuan, J., Smart, P. R., & Shadbolt, N. R. (2007) *Context-Aware Information Retrieval for Enhanced Situation Awareness*. Military Communications Conference (MILCOM), Orlando, Florida, USA.

Bao, J., Tao, J., McGuiness, D. L., & Smart, P. R. (2010) *Context Representation for the Semantic Web*. Web Science Conference, Raleigh, North Carolina, USA.

Bergen, L., Grimes, T., & Potter, D. (2005) How attention partitions itself during simultaneous message presentations. *Human Communication Research*, 31(3), 311-336.

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

Bizer, C., Heath, T., & Berners-Lee, T. (2009) Linked data-the story so far. *International Journal on Semantic Web and Information Systems*, 5(3), 1-22.

Braver, T. S., & Barch, D. M. (2002) A theory of cognitive control, aging cognition, and neuromodulation. *Neuroscience and Biobehavioral Reviews*, 26, 809-817.

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

Carr, N. (2010) *The Shallows: How the Internet Is Changing the Way We Think, Read and Remember*. Atlantic Books, London, UK.

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

Clark, A. (1999) An embodied cognitive science. *Trends in Cognitive Science*, 3(9), 345-351.

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

Clark, A. (2007) Re-Inventing Ourselves: The Plasticity of Embodiment, Sensing, and Mind. *Journal of Medicine and Philosophy*, 32(3), 263-282.

Clark, A. (2008) *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. Oxford University Press, New York, USA.

Clark, A. (2010) Memento's revenge: the extended mind, extended. In R. Menary (Ed.), *The Extended Mind*. MIT Press, Cambridge, Massachusetts, USA.

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

Craik, F. I. M., & Lockhart, R. S. (1972) Levels of processing: a framework for memory research. *Journal of Verbal Learning and Verbal Behaviour*, 11, 671-684.

DeStefano, D., & LeFevre, J. A. (2007) Cognitive load in hypertext reading: A review. *Computers in Human Behavior*, 23(3), 1616-1641.

DeVaul, R., Pentland, A., & Corey, V. (2005) *The memory glasses: subliminal vs. overt memory support with imperfect information*. 7th IEEE International Symposium on Wearable Computers, White Plains, New York, USA.

Diehl, M., & Stroebe, W. (1987) Productivity loss in brainstorming groups: toward the solution of a riddle. *Journal of Personality and Social Psychology*, 53, 497-509.

Dreyfus, H. L. (2009) *On the Internet* (2nd ed.). Routledge, Abingdon, Oxon, UK.

Evans, J. A. (2008) Electronic publication and the narrowing of science and scholarship. *Science*, 321(5887), 395-399.

Eysenck, M. (1974) Age differences in incidental learning. *Developmental Psychology*, 10(6), 936-941.

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.

Hackworth, R. (1952) *Plato's Phaedrus*. Cambridge University Press, Cambridge, UK.

Halpin, H., Clark, A., & Wheeler, M. (2010) *Towards a Philosophy of the Web: Representation, Enaction, Collective Intelligence*. Web Science, Raleigh, North Carolina, USA.

Harnad, S. (1999) The future of scholarly skywriting. In A. Scammell (Ed.), *I in the Sky: Visions of the Information Future*. Aslib, London, UK.

Harnad, S. (2004) Back to the oral tradition through skywriting at the speed of thought. In J.-M. Saluan & C. Vandendorpe (Eds.), *Les défis de la publication sur le Web: Hyperlectures, cybertextes et méta-editions*. Presses de l'ENSSIB, Villeurbanne, France.

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

Hendler, J., 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.

Ho, H., Saeedi, E., Kim, S. S., Shen, T. T., & Parviz, B. A. (2008) *Contact lens with integrated inorganic semiconductor devices*. 21st International Conference on Micro Electro Mechanical Systems, Tuscon, Arizona, USA.

Hurley, D. J., Arbab-Zavar, B., & Nixon, M. S. (2008) The ear as a biometric. In A. K. Jain, P. Flynn & A. A. Ross (Eds.), *Handbook of Biometrics* (pp. 131-150). Springer, Berlin, Germany.

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

Huynh, T. D., Smart, P. R., Braines, D., Shadbolt, N. R., & Sycara, K. (2010) *The Cognitive Virtues of Dynamic Networks*. 4th Annual Conference of the International Technology Alliance (ACITA'10), London, UK.

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.

Kalfoglou, Y., & Schorlemmer, M. (2003) Ontology mapping: the state of the art. *Knowledge Engineering Review*, 18(1), 1-31.

Kirsh, D. (2009) Problem solving and situated cognition. In M. Ayded & P. Robbins (Eds.), *Cambridge Handbook of Situated Cognition*. Cambridge University Press, Cambridge, UK.

Koechlin, E., Ody, C., & Kouneiher, F. (2003) The architecture of cognitive control in the human prefrontal cortex. *Science*, 302(5648), 1181-1185.

Landes, D. (2000) *Revolution in Time: Clocks and the Making of the Modern World*. Viking Press, London, UK.

Lazer, D., & Friedman, A. (2007) The Network Structure of Exploration and Exploitation. *Administrative Science Quarterly*, 52(4), 667-694.

Lee, M. J., & Tedder, M. C. (2003) The effects of three different computer texts on readers' recall: based on working memory capacity. *Computers in Human Behavior*, 19(6), 767-783.

Lingley, A., & Parviz, B. (2008) Multipurpose integrated active contact lenses. *The Neuromorphic Engineer*.

Mason, S., Bashashati, A., Fatourechi, M., Navarro, K., & Birch, G. (2007) A comprehensive survey of brain interface technology designs. *Annals of Biomedical Engineering*, 35(2), 137-169.

Mason, W. A., Jones, A., & Goldstone, R. L. (2005) *Propagation of innovations in networked groups*. 27th Annual Conference of the Cognitive Science Society, Stresa, Italy.

McDonald, S., & Stevenson, R. J. (1996) Disorientation in hypertext: The effects of three text structures on navigation performance. *Applied Ergonomics*, 27(1), 61-68.

Menary, R. (Ed.). (2010) *The Extended Mind*. MIT Press, Cambridge, Massachusetts, USA.

Miller, E., & Cohen, J. (2001) An integrative theory of prefrontal cortex function. *Neuroscience*, 24(1), 167-202.

Newman, S. D., Carpenter, P. A., Varma, S., & Just, M. A. (2003) Frontal and parietal participation in problem solving in the Tower of London: fMRI and computational modeling of planning and high-level perception. *Neuropsychologia*, 41(12), 1668-1682.

Nicolelis, M. (2001) Actions from thoughts. *Nature*, 409(6818), 403-407.

Nixon, M. S., & Carter, J. N. (2006) Automatic recognition by gait. *Proceedings of the IEEE*, 94(11), 2013-2024.

Noë, A. (2004) *Action In Perception*. MIT Press, Massachusetts, USA.

Noë, A. (2009) *Out of Our Heads: Why You Are Not Your Brain, and Other Lessons from the Biology of Consciousness*. Hill & Wang, New York, USA.

Noy, N. F. (2004) Tools for Mapping and Merging Ontologies. In S. Steffen & R. Studer (Eds.), *Handbook on Ontologies*. Springer Verlag, Berlin, Germany.

O'Regan, J. K., & Noë, A. (2001) A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, 24(5), 939-973.

Ophir, E., Nass, C., & Wagner, A. D. (2009) Cognitive control in media multitaskers. *Proceedings of the National Academy of Sciences*, 106(37), 15583-15587.

Pandey, J., Lingley, A., Parviz, B., & Otis, B. (2009) *Toward an active contact lens: Integration of a wireless power harvesting IC*. Biomedical Circuits and Systems Conference (BioCAS'09), Beijing, China.

Pfeifer, R., & Bongard, J. (2007) *How the Body Shapes the Way We Think: A New View of Intelligence*. MIT Press, Cambridge, Massachusetts, USA.

Pfurtscheller, G., Scherer, R., & Neuper, C. (2007) EEG-based brain-computer interface. In R. Parasuraman & M. Rizzo (Eds.), *Neuroergonomics: The Brain at Work*. Oxford University Press, New York, USA.

Poltrock, S., Handel, M., Bowyer, H., & Waggett, P. (2008) *A dynamic model of mission context*. 2nd Annual Conference of the International Technology Alliance (ACITA'08), London, UK.

Rhodes, B. (2003) Using physical context for just-in-time information retrieval. *IEEE Transactions on Computers*, 52(8), 1011-1014.

Rhodes, B., & Maes, P. (2000) Just-in-time information retrieval agents. *IBM Systems Journal*, 39(3-4), 685-704.

Robbins, P., & Ayded, M. (2009) *The Cambridge Handbook of Situated Cognition*. Cambridge University Press, New York, USA.

Roberts, A. C., Robbins, T. W., & Weiskrantz, L. (Eds.) (1998) *The Prefrontal Cortex: Executive and Cognitive Functions*. Oxford University Press, Oxford, UK.

Shadbolt, N. R., & Berners-Lee, T. (2008) Web Science Emerges. *Scientific American*, 299(4), 76-81.

Small, G., & Vorgan, G. (2008) Meet Your iBrain. *Scientific American*, 19(5), 42-49.

Small, G. W., Moody, T. D., Siddarth, P., & Bookheimer, S. Y. (2009) Your brain on Google: patterns of cerebral activation during Internet searching. *American Journal of Geriatric Psychiatry*, 17(2), 116-126.

Smart, P. R. (2008) *Controlled Natural Languages and the Semantic Web.* School of Electronics and Computer Science, University of Southampton, Southampton, England. (Ref: ITA/P12/SemWebCNL)

Smart, P. R. (2010) *Extended Memory, the Extended Mind, and the Nature of Technology-Mediated Memory Enhancement*. 1st ITA Workshop on Network-Enabled Cognition: The Contribution of Social and Technological Networks to Human Cognition, Maryland, USA.

Smart, P. R., Engelbrecht, P., Braines, D., Hendler, J., & Shadbolt, N. R. (2008) *The Extended Mind and Network-Enabled Cognition.* School of Electronics and Computer Science, University of Southampton, Southampton, England. (Ref: ITA/P12/NEC)

Smart, P. R., Engelbrecht, P., Braines, D., Strub, M., & Hendler, J. (2009) *Cognitive Extension and the Web*. Web Science Conference: Society On-Line, Athens, Greece.

Smart, P. R., & Engelbrecht, P. C. (2008) *An Analysis of the Origins of Ontology Mismatches on the Semantic Web*. 16th International Conference on Knowledge Engineering and Knowledge Management (EKAW'08), Acitrezza, Catania, Italy.

Smart, P. R., Engelbrecht, P. C., Braines, D., Strub, M., & Giammanco, C. (2010a) The Network-Extended Mind. In D. Verma (Ed.), *Network Science for Military Coalition Operations: Information Extraction and Interaction*. IGI Global, Hershey, Pennsylvania, USA.

Smart, P. R., Huynh, T. D., Braines, D., & Shadbolt, N. R. (2010b) *Dynamic Networks and Distributed Problem-Solving*. Knowledge Systems for Coalition Operations (KSCO'10), Vancouver, British Columbia, Canada.

Smart, P. R., Huynh, T. D., Braines, D., Sycara, K., & Shadbolt, N. R. (2010c) *Collective Cognition: Exploring the Dynamics of Belief Propagation and Collective Problem Solving in Multi-Agent Systems*. 1st ITA Workshop on Network-Enabled Cognition: The Contribution of Social and Technological Networks to Human Cognition, Maryland, USA.

Stasser, G., & Titus, W. (1985) Pooling of unshared information in group decision making: biased information sampling during group discussion. *Journal of Personality and Social Psychology*, 48, 1467-1478.

Waldron, S., Patrick, J., Morgan, P., & King, S. (2007) Influencing Cognitive Strategy by Manipulating Information Access. *The Computer Journal*, 50(6), 694-702.

Wallas, G. (1926) *The Art of Thought*. Cape, London, UK.

Wilson, R. A., & Clark, A. (2009) Situated Cognition: Letting Nature Take its Course. In M. Ayded & P. Robbins (Eds.), *Cambridge Handbook of Situated Cognition*. Cambridge University Press, Cambridge, UK.

Zhu, E. (1999) Hypermedia Interface Design: The Effects of Number of Links and Granularity of Nodes. *Journal of Educational Multimedia and Hypermedia*, 8(3), 331-358.