Understanding the Cognitive Impact of Emerging Web Technologies: A Research Focus Area for Embodied, Extended and Distributed Approaches to Cognition

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

***Abstract*—Alongside existing research into the social, political and economic impacts of the Web, there is also a need to explore the effects of the Web on our cognitive profile. This is particularly so as the range of interactive opportunities we have with the Web expands under the influence of a range of emerging technologies. Embodied, extended and distributed approaches to cognition are relevant to understanding the potential cognitive impact of these new technologies because of the emphasis they place on extra-neural and extra-corporeal factors in the shaping of our cognitive capabilities at both an individual and collective level. The current paper outlines a number of areas where embodied, extended and distributed approaches to cognition are useful in understanding the impact of emerging Web technologies on future forms of both human and machine intelligence.**

I. Introduction

Over the past couple of years, a debate has emerged concerning the potential impact of the Web on human cognitive function. For the most part, much of the rhetoric in this debate has been negative, with key protagonists, such as Nicholas Carr [1], arguing that the Web is exerting a largely negative effect on our ability to think, read and remember. Instead of enhancing our ability to concentrate, Carr argues, 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 Carr refers 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.

At present, there is a limited amount of scientific evidence to back up many of the claims made by Carr. However, some researchers have suggested that the Web is changing the profile of at least some of our cognitive capabilities. In a recent study on human memory, for instance, Sparrow et al [2] reported that when people are accustomed to using the Web, they are more likely to remember the location of information (or how to retrieve it) than they are to remember the specific details of the information itself. This provides some support for Carr's claims. If we are eschewing biologically-based modes of information storage and recall in favor of Web access and Google search, then we may witness a deterioration in our ability to recall information whenever our access to the Web is compromised.

Other research has revealed that the nature of the hyperlink topologies used within a hypermedia environment may affect our ability to recall information. Studies have shown 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 (e.g., texts in which the links are embedded in the text itself and navigation is not restricted to a linear sequence of nodes) [3]. Again, such results provide some support for Carr's claims about the effects of the Web on our mnenomic capabilities. If our access to factual material is mediated by Web pages replete with links to associated content (as is typically the case in systems such as Wikipedia), we may find ourselves recalling relatively less information than we would have done if we had relied on conventional printed materials.

Whether or not such findings are a cause for alarm may depend on the views we hold with regard to both the Web and the nature of human cognition. In the case of the Web, it is important to remember that the Web is something of a protean beast when it comes to user interaction and information access capabilities. A range of emerging Web technologies, as well as changes in the way we use the Web, all contribute to an ever-changing landscape against which our notions of the cognitive impact of the Web are always likely to be somewhat ephemeral. 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 Web access or usage are likely to exert the same kind of influence on our cognitive profile.

Our view of human cognition also has an effect on how we view research investigating the cognitive effects of Web access. In this respect, approaches that emphasize the embodied, extended and distributed (EED) nature of human cognition have become increasingly popular within the cognitive scientific community [4], [5], [6]. What such approaches have in common is a commitment to the idea that intelligence is often grounded in our ability to factor in the contributions of our extra-neural and extra-corporeal environments. Theoretical and empirical research in embodied cognition thus emphasizes the way in which material embodiment shapes and constrains cognition [6], while proponents of extended cognition see our interaction with aspects of the technological, informational and socio-cultural environment as key to much of our cognitive success [5]. One impact of such views when it comes to understanding the cognitive effects of the Web is to question the significance of research that reveals a change in brain-based cognitive processing. The focus for embodied and extended theories is on the details of our interaction with the external world, and this encourages us to take a systems-level perspective when it comes to issues of cognitive performance. On this view, a new technology could have a beneficial effect on cognitive performance (at both an individual and collective level) even if it seemed to result in an apparent decline in brain-based cognitive function.

In this paper, I will briefly review some of the ways in which emerging Web technologies are of interest to those concerned with EED approaches to cognition. Due to limitations of space, this review will be necessarily brief and therefore highly selective. In Section II, I focus on how the advent of highly portable, wearable devices, such as Google's Project Glass, are poised to transform our traditional notions of Web access. Such technologies situate the Web at the heart of our everyday embodied interactions with the world, and they are, as such, of significant interest to those who emphasize active engagement with the world as part of our cognitive endeavours. Section III raises the possibility that information on the Web may come to constitute part of our personal body of knowledge and beliefs about the world. Inasmuch as this is true, then the epistemic implications of the Web are truly profound, for they suggest that the boundaries of our knowledge are limited only by what the Web makes available to us. Section IV focuses on the use of the Web to store large quantities of personal information and data. A specific point of interest here concerns whether the maintenance of a life-long digital record of one's experiences may alter our opportunity to adaptively regulate our sense of who and what we are. Section V draws attention to the Social Web and describes how aspects of socially-distributed information processing may influence cognitive performance at the collective level. Finally, Section VI describes the impact of the Web on machine intelligence. The main issue here relates to the role played by human agents in advancing the current state-of-the-art in machine-based processing.

II. Embodiment, Cognition and the Web

At first sight, it might appear as though the Web is of little interest to those concerned with EED approaches to cognition. Our predominant vision of online interaction is one in which we are sat in front of a desktop computer, accessing the Web through a conventional browser-based interface (such as Internet Explorer or Google Chrome). In these cases, we are encouraged to view our interaction with the Web as a form of environmentally-decoupled and physically-disembodied interaction. As Canny and Paulos [7] comment: "...cyberspace has been built on Cartesian ideals of a metaphysical separation between mind and body: When we enter cyberspace, even a 3D world, it is the 'mind' that enters...The body stays outside" (pg. 276). This notion of the environmentally-decoupled and physical-disembodied nature of our online interaction contrasts with the main thrust of theoretical and empirical work in the embodied and extended cognition literature. This sees active, real-time engagement with the external world as an important element of our cognitive profile. We might thus be inclined to say that the Web is of little or no interest to proponents of embodied and extended cognition.

One reason to reject this view relates to the changing nature of human interaction with the Web. With the advent of mobile and portable computing solutions, our interactions with the Web are increasingly ones that take place in the context of our everyday sensorimotor engagements with the world. This marks an important shift in the way in which we access the Web, and it is one that enables us to approach the Web as an important part of our environmentally-situated cognitive activities. Emerging technologies are continuing this trend towards a greater interleaving of the Web in our everyday thoughts and actions. The head-mounted augmented reality display device envisioned by Google's Project Glass initiative, for example, promises to make Web information directly available within the visual field. These sorts of innovations situate the Web at the heart of our interactions and engagements with the world, and they make the Web a potent source of interest for those who approach human cognition from an embodied or extended cognition perspective.

III. Extended Knowledge

Devices that increase both the accessibility and perceptual availability of Web-based information have a number of implications for how we view the potential cognitive impact of the Web. One such implication concerns the possibility for Web-based forms of cognitive extension in which the technological and informational elements of the Web come to form part of the supervenience base for (at least some) mental states and processes. As part of their seminal paper on the extended mind, Clark and Chalmers [8] outlined a thought experiment in which a neurologically-impaired individual, Otto, relied on the use of an external resource - a notebook - in order to achieve certain tasks. The main point of the thought experiment was to highlight the similar functional role played by both biological (i.e., the brain/body) and non-biological (e.g., the notebook) resources in supporting at least some cases of intentional action. Inasmuch as the bio-external resources play a role similar to that served by biological resources, Clark and Chalmers claim, we should view their contributions to global behavior as on a functional par. This enables us, at least in some cases, to see bio-external resources as playing a role in the realization of human mental states and processes. When we apply such notions to the Web, we can entertain the possibility of Web-extended minds, or minds in which the technological and informational elements of the Web form part of the physical machinery of the human cognitive system [9].

The possibility of Web-based forms of cognitive extension raises important questions about our future cognitive and epistemic potential. For example, if our access to externally-located information was just as reliably, easily and continuously available as the kind of access afforded by our own bio-memories, then there seems to be no principled reason to suggest that the external information would not count as part of our own personal body of knowledge and dispositional beliefs about the world. This seems to make the boundaries of what we 'know' limited only by the kind of access we have to various sources of environmental information. And if what we have access to is the sum total of human knowledge, as stored on the Web, then the epistemic limits of the Web-extended mind seem to be broadly co-extensive with that of the Web itself!

In order to help clarify all of this, imagine a situation where you are equipped with a mobile networked device (a mobile phone will do) in order to provide wireless access to the Web, an augmented reality head-mounted display device (similar to the technological target envisioned by the Google Project Glass initiative), and a means of controlling information retrieval from the Web in a manner that is sensitive to your ongoing epistemic needs and concerns. Imagine that you use these technologies to retrieve specific items of information from the Web and such items are displayed in your field of view whenever you require them to be available. For the sake of argument, let us say that you are located in an art gallery and you wish to retrieve information about artworks that you have previously encountered. What, we might wonder, would our scientific and social intuitions be about your knowledge of art in such a situation? Would it be appropriate for us to say that you pretty much 'know' everything there is to know about the pieces of art contained in the art gallery, at least in terms of the information that is posted on the Web? If this claim seems profoundly implausible or inappropriate to you, think for a moment about what it is that determines what you think you already know. What seems to determine whether we know or do not know something is not the fact that we are continuously, consciously aware of relevant facts and figures. What seems to count is more the kind of access we have to the relevant information, the fact that when we need to recall the information it is there, easily made available to us by our bio-memory systems. But need our bodies of personal knowledge be so reliant on biologically-based modes of information storage? What if our access to externally-located information was just as reliably, easily and continuously available as the kind of access afforded by our own bio-memories? In this case, it seems, there is no principled reason to suggest that the external information would not count as part of our own body of knowledge and beliefs about the world. As Clark [10] argues:

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

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

In evaluating this kind of technological transformation of our epistemic capabilities, we can discern a number of lines of research. One line of research is focused on the technologies that are used to retrieve and present information from the Web. In all likelihood, such technologies will need to satisfy the kind of conditions that guide our intuitions as to when a form a cognitive extension has taken place. These include conditions associated with the reliability and availability of the device we use to access information, the accessibility of the information to be retrieved and the level of trust we place in the retrieved information [8]. Another factor that is of potential importance concerns the extent to which the technology we use to access Web-based information falls 'transparent in use'[1]. In particular, it has been suggested that the phenomenon of a technology falling transparent in use is closely associated with the emergence of an environmentally-extended cognitive system. Clark [5] thus suggests that the experience of 'seeing through a technology' is a subjective corollary of the technology becoming incorporated into an integrated agent-world cognitive system. Thompson and Stapleton [11] go even further. They argue that "For anything external to the body's boundary to count as part of the cognitive system it must function transparently in the body's sense-making interactions with the environment" (pg. 29). We thus encounter what might be called a 'transparency criterion' in the extended mind debate which serves to guide our sense of when a form of cognitive extension has taken place. In relation to the possibility for Web-based forms of cognitive extension, we can apply the notion of transparent technology to evaluate the various devices we use to access the Web. We might therefore suggest that in order for some form of agent-Web coupling to count as a genuine case of cognitive extension, the devices used to interact with the user should be transparent in use. This not only serves as a test for candidate cases of cognitive extension, it also provides the basis for design and engineering

[1]The notion of a technology falling 'transparent in use' is used to characterize the subjective shift in experience that is associated with the automatic and fluid use of a device (classic examples include the blind man's cane or the carpenter's hammer). Once a technology falls transparent in use, the user is no longer aware of all the details of technology use but is rather focused on the task at hand.

efforts that seek to bring Web-extended minds into existence.

The way in which information is represented on the Web is also relevant to the emergence of Web-based forms of cognitive extension. Of particular relevance here is the transition to what has been called the Web of Data [12], which sees the Web as a globally-distributed database in which data linkages are established by dereferenceable Uniform Resource Identifiers (URIs). As I have discussed elsewhere [9], I believe this transition to a Web of Data plays an important role in enabling the Web to function as a component of extended cognitive systems. One reason for this concerns the accessibility of specific items of information - the fact that is possible to retrieve isolated pieces of information in a wide variety of different contexts. Another reason relates to the fact that the information content is much more amenable to machine-based processes that can find, filter and format data in ways that are optimally suited to human end-users' specific information needs and concerns. Finally, the move away from conventional Web pages (encoded using HTML) to linked data formats (the data format encountered on the Web of Data) opens up a range of presentational capabilities that can be used to guide thought and action in particular ways. Thus, when we think of the potential of new devices, such as Google's Project Glass and its successors, we should not necessarily think of their presentational capabilities as being limited to the display of conventional (HTML) Web pages. Instead, we should think of a whole variety of different data-driven presentational capabilities, some assuming the form of simple natural language statements and instructions, others relying on the use of graphical cues and prompts. In addition, the notion of augmented, mixed or blended reality enables us to think of Web-based information being used to create virtual overlays on the physical environment, enriching the range of affordances to which our sensorimotor systems are attuned. In general, 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 use of linked data formats within the context of the Web of Data marks an important step towards the realization of these kinds of capabilities.

A third area of research interest concerns the factors that govern our meta-cognitive judgements regarding when we feel we know something. This is relevant to the case outlined earlier, where we entertained the possibility that sustained exposure to a reliable information resource could lead to subjective shifts in our sense of what we felt we knew about particular parts of the world. Within the psychological literature, our sense as to what we feel we know has been referred to as the 'feeling of knowing' (FOK). According to one psychological model, such feelings are grounded in the products of the retrieval process itself [13]. Thus, when we attempt to retrieve some item of information from bio-memory, our FOK is based on the amount of information that gets retrieved. This need not be the actual target item itself; it could be a range of fragmentary material, such as "semantic attributes, episodic information, and a variety of subtle activations emanating from other sources" (pg. 159) [13]. A key question for future research is thus the extent to which similar processes, incorporating bio-external information resources, could support feelings similar to those encountered in the case of bio-memory. If our retrieval attempts against the Web were rewarded with the immediate presentation of target items, or related material, would we experience the same FOK as encountered when relying on brain-based forms of information recall? Such questions obviously impact on technology design; for, inasmuch as the FOK is possible with forms of extended memory, then the question for engineers is whether Web-enabled technologies can (ever?) meet the requirements (in terms of retrieval speed, presentational format, and retrieval initiation mechanisms) that cause subjective feelings of knowing to arise.

IV. Memories for Life

In Section III we saw that questions concerning the impact of the Web on our epistemic capabilities typically focused on issues of semantic memory - the component of our memory that deals with the storage and recall of particular facts and conceptual knowledge. Other uses of the Web, brought about by emerging technologies, may be poised to influence other aspects of our memory as well. Of particular interest in this respect is the use of the Web to store large quantities of personal data tracking our life experiences, projects, goals, interests, social acquaintances, health status, and so on. Such data (in the form of photos, videos, and online textual postings) is already being stored as a result of our participation in social media and social networking sites, and the use of cloud computing resources and lifelogging technologies is likely to lead towards an even greater tendency to maintain traces of our past in the online world.

The cognitive impact of this mode of Web use, characterized by the online storage of large quantities of personal data, is still largely unknown. One issue for research in this area relates to the potential impact of Web-based storage on our autobiographical memories[2]. It is known that technologies that record a trace of our daily activities and events can enhance our memory of the past. The use of wearable cameras, such as SenseCam, has been a particular focus of research attention in this area, with studies showing that SenseCam images can work to support the retrieval of autobiographical memories in clinical populations [15], [16]. There are also efforts underway to explore the potential of lifelogging technologies and semantic modes of information representation to support the compilation of autobiographical knowledge bases and enhance autobiographical memory [17]. The Memories for Life initiative (http://www.memoriesforlife.org/), for example, is a multidisciplinary research effort that aims to advance our understanding of memory and develop technologies that

[2]Autobiographical memory refers to the memories we have of specific events across the lifespan, and it includes the body of personal knowledge (autobiographical knowledge) that is associated with our life experiences [14].

enhance our mnenomic capabilities. The results of such studies and research efforts are important because they shed light on the way in which online records of our daily activities and life events are able to enhance our autobiographical memories. Based on the value of at least some forms of lifelogging data to serve as a retrieval cue, we can expect the move towards the online storage of personal data to improve our ability to accurately recall events and episodes from our past life.

All this leads to the conclusion that the Web may play a largely positive role in supporting aspects of human memory. In the case of semantic memory (see Section III) we encountered the possibility that specific forms of interaction with the Web could enhance our ability to recall facts and information relating to particular domains of interest, and in the case of autobiographical memory we have seen how our ability to use the Web to store details of our daily activities could enhance our ability to recall past life events. In spite of this, however, when we think of memory enhancement and performance, we should be mindful about the various functions that different forms of memory play in our lives. The ability to accurately and consistently recall previously encountered information is clearly important to some forms of memory (e.g., semantic memory); however, it is not clear that accuracy and indelibility are always a target for memory enhancement. Autobiographical memories, for example, are believed to play a important role in determining our sense of who and what we are (i.e., our personal identity) [18], As a result, an improved ability to recall autobiographical memories may affect the extent to which we can change our identities across the lifespan. As Conway [19] points out, our autobiographical memories contribute to our knowledge about who we are, and this in turn "constrains what the self is, has been and can be" (pg. 594). All this calls into question the value of maintaining an accurate, indelible trace of all aspects of our past experience; for a certain degree of malleability and effaceability in our autobiographical memories may be essential to our capacity to view ourselves and lives in a positive way - in particular, in a way that promotes psychological well-being. It has thus been suggested that one of the functions of autobiographical memory is to maintain a largely favorable view of the self [20], [18], and mnenomic distortion may work to support this function. Wilson and Ross [18] argue that people are motivated to push negative events further back in order to maintain a favorable view of themselves in the present. They show how people can distance themselves from negative events by pushing them into the distant past, so as to make them no longer relevant to the current self's well-being. In addition, positive events can be pulled forward in time so that the current self can continue to take credit for past successes. The function of these distortions, according to Wilson and Ross [18], is to enable people to create and maintain a coherent - and largely positive - view of their present selves and associated circumstances. There are, of course, times when accuracy in recall is important. The recall of specific facts from semantic memory seems to be a case in point. However, we should perhaps be cautious of efforts that strive to provide blanket improvements in mnemonic accuracy across all types of memory. If accuracy is not at the functional heart of a particular kind of memory (e.g., autobiographical memory), then it is difficult to see how technological interventions that focus on accurate recall can really contribute to an improvement or enhancement of that particular aspect of mnemonic functioning. In general, a technology improves a cognitive capability if it delivers a profile of enhanced performance in the same areas as that targeted by the original capability. If the technology works to improve performance in some other area, then the technology is not so much improving the existing capability as installing a new kind of capability, with all the attendant risks that that entails. If the primary function of our autobiographical memory system is to support self and social functions (e.g., maintaining a positive view of the self, bolstering self-esteem, creating a coherent life story, and providing material to stimulate conversation), then it is unclear whether Web technologies that improve the accuracy of recall can actually serve to enhance autobiographical memory. Inasmuch as such technologies work in opposition to existing processes, they may in fact do more harm than good. Rather than focus on accurate recall, a much better focus of attention for such technologies might be to concentrate on how to support the human individual in constructing renditions of the past that enable them to function better in both the psychological and social domains. In some cases at least, this might mean that the best thing memory technologies could do would be to help us modify our recollections of the past...or even dispense with them altogether.

V. Collective Cognition and the Social Web

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 social interaction and coordinating collective efforts. Ever since the advent of Web 2.0, which is characterized by greater levels of user participation in the creation, maintenance and editing of online content, the Web has provided ample opportunities to support various forms of socially-distributed information processing. In addition, the recent surge in social media sites (e.g., YouTube), social networking systems (e.g., Facebook) and microblogging services (e.g., Twitter) has opened up new ways for people to interact, communicate and share information content. We are increasingly seeing the emergence of what we might call the 'Social Web': a suite of applications, services, technologies, formats, protocols and other resources, all united in their attempt to both foster and support social interaction.

Perhaps unsurprisingly, the Social Web is of considerable interest to those who adopt EED approaches to cognition. This is because human cognition is often seen as a socioculturally situated activity, and great emphasis is placed on the role of social forces and factors in shaping our cognitive capabilities. One point of interest here concerns the way in which some cognitive processes, such as reasoning, remembering and problem-solving, might be seen as socially distributed [21], [4]. Within the context of the Web and Internet Science community, the advent of the Social Web has given rise to an increasing interest in the socially-distributed nature of human cognition [22], [23], [24], [25], and this interest has been accentuated with the recent explosion in social computing [26], human computation [27] and collective intelligence [28] systems. Such technologies focus attention on the ways in which the Web may be used to exploit the latent 'socio-cognitive capital' possessed by large numbers of geographically dispersed individuals.

Because of the kinds of opportunities it affords for large-scale collaboration, information sharing, and the coordination of collective efforts, the Web emerges as a seemingly natural platform to realize advanced forms of collective intelligence. However, in spite of the apparent potential of the Web to support socially-distributed cognition, it is important to understand that not all forms of Web-based social interaction and information exchange necessarily lead to improvements in collective cognitive processing. It is known, for example, that the rate at which information and ideas are distributed through a social network can have a profound effect on group-level cognitive outcomes, and this highlights a source of tension in our attempts to engineer systems that support socially-distributed cognition in Web-based contexts. On the one hand, we are usually inclined to countenance high-bandwidth communication systems that feature high levels of connectivity and which maximize the efficient and widespread dissemination of information to all members of a community. On the other hand, we encounter a range of findings in the social psychological and multi-agent simulation literature that suggest that such systems may not always deliver the best outcomes in terms of collective cognitive performance. In some situations, at least, the rapid communication of information and ideas does not always serve the collective cognitive good: precipitant forms of information sharing can sometimes subvert rather than support socially-distributed cognition [21], [29], [30].

What all this means, of course, in terms of our attempt to support socially-distributed cognition on the Web, is that we need to develop a better understanding of the effect that different forms of information flow and influence have on collective cognitive outcomes. One factor that has emerged as an important focus of research attention, in this respect, is the structure of the communication network in which individuals are embedded. Research has shown that the structure of the communication network shapes the flow of information between individuals, and this can lead to different effects on group-level performance [25]. Another factor that has proved of considerable research interest concerns the amount of feedback that is given to individuals about the progress or status of collective cognitive processing. Lorenz et al [31], for example, have shown that by providing feedback about the judgements of others, performance in a collective estimation task is undermined. Subjects provided with high levels of feedback settled on estimates that were, at the collective level, worse than those seen in situations in which subjects received no feedback at all. In accounting for these results, Lorenz et al [31] posit a 'social influence effect' in which the feedback about other user's ratings is deemed to progressively reduce the diversity of ratings within the group without a corresponding improvement in group-level accuracy. These results suggest that although the Web provides an environment in which a variety of kinds of information can be gathered during the course of socially-distributed information processing, not all of this information should be made available to the individual agents engaged in the process. Instead, the results call for a more nuanced approach in which the system works to adaptively regulate the availability of different kinds of information in ways that are sensitive to the nature of the task that is being performed, as well as the psychological propensities of the participating agents. In essence, what is required is a way of dynamically organizing the setup of Web-based socio-technical systems in order to meliorate group-level cognitive processes in a variety of different task contexts.

VI. Human-Assisted Machine Intelligence

When we think of the cognitive implications of the Web in an individual or collective context, our attention typically tends to focus on the ways in which Web technologies can enhance or augment the cognitive performance of individual human agents, or collections thereof. In this case, we see the Web from the perspective of human agents who are interested in pressing maximal cognitive benefit from whatever representational and computational resources the Web has to offer, and we ask questions concerning how the Web can help us in achieving certain cognitive outcomes. There is, however, another perspective we can take with respect to human participation on the Web. This perspective focuses on human agents as the locus of particular kinds of capabilities that subtend the epistemic, cognitive, perceptual, behavioral, social, moral, emotional and affective domains, and it invites us to ask questions concerning how these capabilities can be used to support various forms of machine-based processing.

In order to help us see things from this alternative perspective, consider the use of CAPTCHAs on the Web. CAPTCHAs are questions that are posed by a computer in order to check that the user of a system is a human agent rather than a machine entity. They are typically encountered on the Web in the form of distorted word images that are used to prevent automated access to some resource or service. The human end-user, when confronted with the CAPTCHA, must identify the distorted word and respond by entering the correct form of the word in a text field. This is something that human agents find relatively easy to do because of their capacity for pattern recognition/completion and their familiarity with lexical representations. It is, however, something that conventional computers still find difficult to accomplish. Given the aim of distinguishing human from non-human users, of course, CAPTCHAs tend to focus on the range of capabilities that distinguish human from machine intelligence (or at least the current state-of-the-art in machine intelligence). Aside from the obvious ability of humans to identify distorted word forms, CAPTCHAs could also thus target humans' abilities for aesthetic evaluations, moral judgements, the processing of linguaform representations and the identification of objects from multimedia resources. All of these capabilities are ones that we are able to realize by virtue of our particular form of biologically-based and socially-situated intelligence, and they are ones that often rely on the fact that we are embodied agents that are capable of sensing and acting in the real world. These capabilities are obviously ones that are of interest to those working in the area of artificial intelligence; however, they are ones that machines have, at least thus far, failed to fully emulate.

Enter reCAPTCHA [32]. reCAPTCHA is a system that capitalizes on the responses humans make to CAPTCHAs in order to support various forms of machine-based processing. One use of human responses to the aforementioned distorted word images, for example, is to support the automatic digitization of archived texts that are too degraded or corrupted to be properly handled by optical character recognition (OCR) technology. What we see here is a situation in which the information processing capabilities of the Web's human user community is being exploited in order to help realize a task assigned to a computerized system. This is more-or-less the reverse of the situation we encountered earlier in the paper where the emphasis was on the exploitation of computerized resources in order to help human agents achieve particular tasks (e.g., factual recall). By taking an EED approach to cognition, therefore, we are able to see the socio-technical environment of the Web as a network of biological, technological and informational resources, each of which is capable of being assimilated into bio-technologically hybrid information processing routines. When our concern is with human intelligence, our attention focuses on the representational and computational resources of the Web, and we ask how these resources can be used to extend the reach of our cognitive capabilities. However, when our concern is with machine intelligence, our attention shifts to the social environment in which the Web is embedded and we ask how the capabilities of human agents can be used to advance the state-of-the-art in machine-based processing. In both cases, what matters is the way in which a rather diverse (albeit complementary) range of representational and computational resources are being co-opted into information processing routines that yield intelligent outcomes. One virtue of the Web, in this respect, is that it enables us to tap into the capabilities of human agents in a manner (and on a scale) that has never been seen before. The unique feature of the Web is the fact that it is a social technology that interfaces with a large proportion of humanity. By firmly embedding itself within a human social environment, the Web opens up a range of opportunities to *incorporate*[3] human agents into episodes of machine-based processing. We can therefore approach the Web in the same way that EED cognition theorists often approach the human agent: as a system that is capable of factoring in the contributions of a wider set of resources in order to achieve forms of intelligent processing that would otherwise be difficult or impossible to realize.

[3]The notion of incorporation here is analogous to that proposed by extended mind theorists who often see cognitive extension as involving the active assimilation of bio-external resources into a larger bio-technologically hybrid information processing economy [5].

VII. Conclusion

Alongside existing research into the social, political and economic impacts of the Web, there is also a need to explore the effects of the Web on our cognitive profile. EED approaches to cognition are important here because of the emphasis they place on extra-neural and extra-corporeal factors in shaping the course of our cognitive endeavours. As the range of our interactive opportunities with the Web expands under the influence of emerging technologies, such as wearable devices, mixed/blended reality solutions and the 'Web of Things', so approaches that emphasize the details of material embodiment and physical embedding to our cognitive performance are likely to become of ever greater significance.

One of the advantages of EED approaches to cognition is that they open up new lines of enquiry concerning the effect of the Web on our cognitive capabilities. In the case of memory, for example, we saw that issues of cognitive extension motivate a consideration as to whether the Web can serve as a form of extended memory. It was suggested that under certain conditions, the Web might be able to serve as an external repository of knowledge and beliefs, and, inasmuch as this is true, we might expect the Web to lead to a significant improvement in our ability to recall factual information.

While performance enhancement in the case of semantic memory clearly relies on our ability to accurately recall large bodies of factual information, we saw that enhancement in the case of autobiographical memory may be linked to a rather different set of performance metrics. Enhanced performance, in this case, may not be so much about the ability to recall information as the ability to manage one's access to the past in a way that supports one's current and future well-being. Perhaps this sounds a cautionary note for all attempts to enhance cognitive performance. In all cases, our attempts at performance enhancement should consider the way in which cognitive processes relate to aspects of our physical, social and psychological well-being.

In addition to its effects on aspects of individual cognitive functioning, the Web also plays a role in supporting aspects of socially-distributed cognition. There is a significant body of work to be undertaken in terms of improving our understanding of the factors that influence cognitive success in these situations. We have seen that the structure of social networks can play a role in determining performance outcomes, and these effects may be specific to particular tasks [25]. We have also seen that feedback regarding the views and judgements of others can sometimes work to undermine collective intelligence capabilities [31]. Such results serve to remind us that not all the features of the Web (e.g., its capacity for rapid information dissemination) may work to enhance cognitive performance at the collective level. Further research is required to understand how the Web can be exploited to the collective cognitive benefit of those who use it.

Finally, EED approaches to cognition help us to adopt a different perspective when it comes to the Web's role in supporting intelligent processing. The notion of cognitive extension encourages us to see biotechnological hybridization as a key feature of human intelligence [10], and, in this respect, the increasing penetration of the Web into every aspect of our lives opens up a range of opportunities for the formation of cognitively-potent biotechnological mergers. However, we can also see the Web as a means of incorporating human agents in machine-based information processing, yielding new forms of capabilities that advance the current state-of-the-art in machine intelligence. Just as human cognitive performance is seldom one that relies solely on the representational and computational resources of the biological brain, so too advanced forms of machine intelligence may be ones that benefit from a degree of biotechnological hybridization. In this respect, the increasing penetration of the Web into every aspect of our lives makes it perfectly poised to participate in the realization of socially-extended forms of machine intelligence: ones in which the computational and representational resources of society at large are used to extend the current reach of machine-based capabilities.

Acknowledgment

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

References

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

[2] B. Sparrow, J. Liu, and D. Wegner, "Google effects on memory: Cognitive consequences of having information at our fingertips," *Science*, vol. 333, no. 6043, pp. 776–778, 2011.

[3] D. DeStefano and J. A. LeFevre, "Cognitive load in hypertext reading: A review," *Computers in Human Behavior*, vol. 23, no. 3, pp. 1616–1641, 2007.

[4] E. Hutchins, *Cognition in the Wild*. Cambridge, Massachusetts, USA: MIT Press, 1995.

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

[6] L. A. Shapiro, *Embodied Cognition*. Abingdon, Oxfordshire, UK: Routledge, 2011.

[7] J. Canny and E. Paulos, "Tele-embodiment and shattered presence: Reconstructing the body for online interaction," in *The Robot in the Garden: Telerobotics and Telepistemology in the Age of the Internet*, K. Goldberg, Ed. Cambridge, MA, USA: MIT Press, 2000, pp. 276–294.

[8] A. Clark and D. Chalmers, "The extended mind," *Analysis*, vol. 58, no. 1, pp. 7–19, 1998.

[9] P. R. Smart, "The web-extended mind," *Metaphilosophy*, vol. 43, no. 4, pp. 426–445, 2012.

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

[11] E. Thompson and M. Stapleton, "Making sense of sense-making: Reflections on enactive and extended mind theories," *Topoi*, vol. 28, no. 1, pp. 23–30, 2009.

[12] T. Heath and C. Bizer, *Linked Data: Evolving the Web into a Global Data Space*, 1st ed., ser. Synthesis Lectures on the Semantic Web: Theory and Technology. Morgan and Claypool, 2011.

[13] A. Koriat, "The feeling of knowing: Some metatheoretical implications for consciousness and control," *Consciousness and Cognition*, vol. 9, no. 2, pp. 149–171, 2000.

[14] A. Baddeley, M. W. Eysenck, and M. C. Anderson, *Memory*. Hove, East Sussex, UK: Psychology Press, 2009.

[15] A. R. Doherty, S. E. Hodges, A. C. King, A. F. Smeaton, E. Berry, C. Moulin, S. Lindley, P. Kelly, and C. Foster, "Wearable cameras in health: The state of the art and future possibilities," *American Journal of Preventive Medicine*, vol. 44, no. 3, pp. 320–323, 2013.

[16] R. Brindley, A. Bateman, and F. Gracey, "Exploration of the use of SenseCam to support autobiographical memory retrieval within a cognitive behavioural therapeutic intervention following acquired brain injury," *Memory*, vol. 19, no. 7, pp. 705–712, 2011.

[17] K. O'Hara, M. Tuffield, and N. Shadbolt, "Lifelogging: Privacy and empowerment with memories for life," *Identity in the Information Society*, vol. 1, no. 1, pp. 155–172, 2009.

[18] A. Wilson and M. Ross, "The identity function of autobiographical memory: Time is on our side," *Memory*, vol. 11, no. 2, pp. 137–149, 2003.

[19] M. Conway, "Memory and the self," *Journal of Memory and Language*, vol. 53, no. 4, pp. 594–628, 2005.

[20] S. Bluck, "Autobiographical memory: Exploring its functions in everyday life," *Memory*, vol. 11, no. 2, pp. 113–123, 2003.

[21] E. Hutchins, "The social organization of distributed cognition," in *Perspectives on Socially Shared Cognition*, L. Resnick, J. Levine, and S. Teasley, Eds. Washington DC, USA: The American Psychological Association, 1991.

[22] E. H. Chi, "The social web: Research and opportunities," *Computer*, vol. 41, no. 9, pp. 88–91, 2008.

[23] E. H. Chi, P. Pirolli, B. Suh, A. Kittur, B. Pendleton, and T. Mytkowicz, "Augmented social cognition," in *AAAI Spring Symposium on Social Information Processing*, California, USA, 2008.

[24] E. H. Chi, "Augmented social cognition: using social web technology to enhance the ability of groups to remember, think, and reason," in *35th SIGMOD International Conference on Management of Data*, Providence, Rhode Island, USA, 2009.

[25] M. Kearns, "Experiments in social computation," *Communications of the ACM*, vol. 55, no. 10, pp. 56–67, 2012.

[26] M. Parameswaran and A. B. Whinston, "Research issues in social computing," *Journal of the Association for Information Systems*, vol. 8, no. 6, pp. 336–350, 2007.

[27] A. Quinn and B. Bederson, "Human computation: a survey and taxonomy of a growing field," in *Annual Conference on Human Factors in Computing Systems (CHI'11)*, Vancouver, British Columbia, Canada, 2011.

[28] E. Bonabeau, "Decisions 2.0: The power of collective intelligence," *MIT Sloan Management Review*, vol. 50, no. 2, pp. 45–52, 2009.

[29] D. Lazer and A. Friedman, "The network structure of exploration and exploitation," *Administrative Science Quarterly*, vol. 52, no. 4, pp. 667–694, 2007.

[30] P. R. Smart, T. D. Huynh, D. Braines, and N. R. Shadbolt, "Dynamic networks and distributed problem-solving," in *Knowledge Systems for Coalition Operations (KSCO'10)*, Vancouver, British Columbia, Canada, 2010.

[31] J. Lorenz, H. Rauhut, F. Schweitzer, and D. Helbing, "How social influence can undermine the wisdom of crowd effect," *Proceedings of the National Academy of Sciences*, vol. 108, no. 22, pp. 9020–9025, 2011.

[32] L. von Ahn, B. Maurer, C. McMillen, D. Abraham, and M. Blum, "reCAPTCHA: Human-based character recognition via web security measures," *Science*, vol. 321, no. 5895, pp. 1465–1468, 2008.