Towards a Philosophy of the Web: Representation, Enaction, Collective Intelligence

Harry Halpin
School of Informatics
University of Edinburgh
Informatics Forum
10 Crichton Street, Edinburgh
EH8 9AB United Kingdom
hhalpin@inf.ed.ac.uk

Andy Clark
Department of Philosophy
University of Edinburgh
Dugald Stewart Building
3 Charles Street, Edinburgh
EH8 9AD United Kingdom
aclark3@staffmail.ed.ac.uk

Michael Wheeler
Department of Philosophy
University of Stirling
Stirling
FK9 4LA United Kingdom
m.w.wheeler@stir.ac.uk

ABSTRACT

We present some initial forays into the questions that underlie the philosophy of the Web around the notions of representation, enactive search, the extended mind, and collective intelligence.

Keywords

philosophy, representations, enaction, collective intelligence

There is an emerging vision of the human mind as essentially a social organ apt to make extensive and transformative use of whatever forms of local and global scaffolding other agents and technologies provide. In an increasingly wired and networked world, our very nature as cognitive beings is gradually changing.

The World Wide Web is a remarkable triumph of incremental computational engineering. In the wake of its technological success and the structural change it has effected on human social organization, Web designers and researchers are being forced to confront a range of foundational issues with clear philosophical dimensions. These include old philosophical issues in modern guises issues concerning knowledge, identity and trust as well as new questions raised by the increasingly complex ways in which the Web is embedded in the larger world. Such new questions concern, for example, the character and status of web objects such as websites and mash-ups, the understanding of authorship within new collaborative and collective creative ensembles, and the relation between, on the one hand, the structure and functioning of the web and, on the other, the strengths and weaknesses of basic biological cognition.

1. IS PHILOSOPHY PART OF WEB SCIENCE?

Sir Tim Berners-Lee, the British inventor of the Web, has challenged philosophy to contribute to the future construction of the Web, calling the architects behind the Web 'philosophical engineers" [3]. Our hypothesis is that profitable philosophical engagement with the Web will be achieved through the lens of contemporary debates in philosophy of mind and cognitive science over what is sometimes called *4E (embodied, embedded, enactive, extended) cognition*. Indeed, it is this area of philosophy that has attracted the most interest from leading thinkers about the Web, although to date there are very few examples of philosophers of 4E cognition engaging directly with the Web. As we explain below, a number of critical and fundamental questions that Web researchers are beginning to confront are intertwined with issues at the forefront of recent work on 4E cognition. Importantly, however, the intellectual traffic here is not one-way. A carefully crafted philosophy of the Web will not merely draw on the aforementioned philosophical debates, it will make important contributions to them. The study of human-Web couplings provides a powerful way to pursue several unresolved and controversial contemporary philosophical issues. In this way, practical interests in the design and use of web-based technologies dovetail with foundational questions concerning the nature of minds and persons.

This new philosophy of the Web should go beyond much 'new media studies' work on the Web by engaging directly with certain pressing scientific and engineering concerns faced by Web architects. What was missing from this earlier work was the productive engagement with the relevant philosophical community that the present project will foster. However, much of the fault of this lies on the shoulders of philosophy. The debate over relevance-sensitivity to be found in the 4E philosophical literature has seldom been connected to the performance of search engines or the Web more generally [9]. Similar comments might be made about enactivist thinking. By contrast, the extended mind hypothesis is often accepted as almost intuitively obvious in Web circles, even though, in philosophical circles, it is the subject of ongoing critical debate [1]. The established interest in the extended mind by the Web community is indicated by, for example, work of computer scientists and psychologists on applying the extended mind hypothesis to the Web [17]). Moreover, the extended mind hypothesis was advanced as a possible foundation for a philosophy of the Web in a 2008 exchange between Halpin [11] and Wheeler [20] in the *American Philosophical Association Newsletter on Philosophy and Computers*. Based upon this previous work, we can outline three linked themes that a philosophy of the Web that is based on empirical work and cognitive science should address: the problem of representations on the Web, enactive search, and collective intelligence.

Web Science Conf. 2010, April 26-27, 2010, Raleigh, NC, USA.

2. REPRESENTATIONS AND THE WEB

Cognitive science and artificial intelligence (AI) have traditionally appealed to internal and neurally-realised representations to explain intelligence. But exactly how the notion of representation should be understood, and the extent to which neurally located representations are necessary for intelligence, have long been vexed issues, especially in the philosophical literature. The Web provides a new impetus for investigation into the notion of representation, because the advent of the Web has seen an explosion of novel external representations (e.g. hypertext web-pages) that, through complex, iterated interactions with human users, enable intelligent information retrieval, complex commercial activity, and social communication and co-ordination. We shall explore the relationship between these two seemingly different representational contributions to intelligent activity, asking whether there is a single account of representation that applies to both of them.

These are theoretical questions with practical implications. The Semantic Web (also known as Linked Data) is a project launched by Berners-Lee et al. [5] and deployed by august bodies including the UK government (e.g. in the recent initiative to release government data, `http://data.gov.org.uk`) and the BBC. Its goal is to build large-scale structured knowledge representation systems using the Web. As such, it has been identified by some computer scientists as an attempt to restart the project of classical artificial intelligence [10]. If this is right, the Semantic Web will ultimately face the problems that, according to some critiques, plagued its intellectual predecessor [6]. One such problem is the frame problem, the recalcitrant difficulty of determining, in a wholly mechanistic manner, which items of information from a huge memory store are relevant (and which are irrelevant), and how retrieved items of information should be updated, within and across changing contexts of activity [15]. Influenced by phenomenological philosophers such as Heidegger (who stressed embeddedness) and Merleau-Ponty (who additionally stressed embodiment), thinkers such as Dreyfus have argued that the root cause of such problems is the assumption that the mechanistic processes at work are representation-guided [8]. Perhaps anti-representationalist embodied cognition has itself been taken too far? Alternative remedies that depend on a reconceived, Web-friendly notion of representation, and identify the implications for how the Web may meet the challenge of relevance-sensitivity.

3. ENACTIVE SEARCH

Currently, access to representations on the Web is mediated through search engines like Google. One of the keys to the practical success of such search engines is that they use massive amounts of statistics gathered from user-actions and user-choices on the Web to dynamically adapt their algorithms to find appropriate content, and thus to grapple with the issues of relevance and context highlighted by the frame problem. The more information a search algorithm has about the conduct and interests of users, the better its adaptation, a trend accelerated by 'Web 2.0' technologies such as social networking and collaborative tagging.

We hypothesize that the complex adaptive dynamics of such statistics-driven user-action-based search may be illuminated using the philosophical concept of enaction (roughly, the idea of 'laying down a path in walking,' i.e., of actions that change the world in ways that feed and structure those very actions, either now or at some future time [18]. Although much of the existing work on enaction has focused on the biological individual, an enactive paradigm can also be applied to the collective effects of our use and navigation of the Web. From this theoretical standpoint we shall investigate how search engines highlight a number of key issues in the way we think about minds, persons and collective endeavours. These range from a kind of 'quantification' of Wittgenstein's maxim 'meaning is use' as a form of statistical language processing (an intellectual legacy dating back to the early search engines produced in Cambridge in the 1960s [21]) to the empowering and disturbing (in about equal measure) vision of the not-too-distant future, described recently by Google's CEO Eric Schmidt, in which Google is connected 'straight to your brain'.[1]

4. COGNITIVE EXTENSION AND COLLECTIVE INTELLIGENCE

According to the extended mind hypothesis as presented by Clark and Chalmers [7], cognitive processes are not always confined within the boundaries of skin and skull. Under certain conditions in which biological brains and bio-external scaffoldings work together as integrated processing ensembles, cognition may extend into the world. The nature of the Web opens up this controversial philosophical issue in a distinctive and distinctively problematic way. Perhaps external representations on the Web, when integrated appropriately into the processes that govern an agents behaviour, may count as parts of that agents cognitive architecture. But now assume that multiple individuals are able to access the same external representation, such as a Google Map, and that they can update it in near-real time. Here, it seems, more than one person may deserve cognitive credit for, and have cognitive ownership of, a representation that augments their own individual intelligence [11]. What is still up for grabs is exactly which bio-technological couplings yield genuine extensions of cognition, rather than merely novel supporting environmental structures in which internally-constituted cognition may function.

Finally, an analysis of the social dimension of Web-enabled distributed intelligence will help us to understand our increasing dependence on the Web, not only for information gathering but also for socially co-ordinating action (via tools like social networking and microblogging). An adequate account of such collective intelligence (or the intelligence of collectives) on the Web will revisit existing theories of distributed cognition in the light of the case for the Web-extended mind [13]. A central challenge will be to analyse the conditions under which users trust, by responding unreflectively and uncritically to, the collectively maintained information retrieved from the Web. Plausibly (cf. information retrieved from internal memory), such trust is a necessary condition for the external representations involved to qualify as cognitive extensions. From Wikipedia to Google, this ability to trust information on the Web is one of the most pressing problems facing the Web today [16], so a detailed philosophical analysis of this topic promises to have a significant impact on the practice of Web engineers designing these systems.

[1]See `http://techcrunch.com/2009/09/03/google-ceo-eric-schmidt-or`

5. FROM THE EXTENDED MIND TO THE WEB

In order to hone some of these points, we put forward a speculative extension to the Extended Mind hypothesis that shows how the Web may effect the argument. As put forward by Wheeler, "online intelligence is generated through complex causal interaction in an extended brain-body-environment system" [19]. We can press on this point make room for an active role of representations in general, and for the Web in particular, "an active externalism, based on the active role of the environment in driving cognitive processes" [7]. For example, a representation can be stored in the memory "inside" the head of an agent in some neural state, but it can just as easily be stored outside in a map. The debate over the existence of internal representations is an empirical debate best left to neuro-scientific work. However, what is less up for debate seems to be that representations at least exist externally for particular agents. Finding those representational neural states are difficult, but let us not deny the existence of maps!

In the original argument for the Extended Mind Hypothesis, Clark and Chalmers introduce us to Otto, a man with an impaired memory who navigates about his life via the use of maps and other such notes in his notebook [7]. Otto wants to navigate to the Metropolitan Museum of Modern Art in New York City from his house in Brooklyn, but to do with his impaired memory he needs a map.[2] In order to arrive at the museum, Otto needs a map whose components are in some correspondence with the world he must navigate in order to get to the museum. Otto can find a map to the Museum of Modern Art that exists for the precise purpose of navigating individuals to the museum. It is hard to deny that a map is representational in the sense we have presented above, as it is a representation whose target is the various streets on the way to the Museum. The map is just an external representation in the environment of Otto, and can drive the cognitive processes of an Otto in a similar fashion to the way that classical AI assumed internal representations in Otto's head did. Clark and Chalmers point out that if external factors are driving the process, then they deserve some of the credit: "If, as we confront some task, a part of the world functions as a process which, were it done in the head, we would have no hesitation in recognizing as part of the cognitive process, then that part of the world is (so we claim) part of the cognitive process" [7].

Humans have difficulty maintaining any sort of coherent memory in front of the vagarities of their dynamic environment. Technology allows this weakness to be turned into a strength, for the consequence is that humans who have a way to maintain an external memory capable of holding representations, even if such a memory is technically outside of their biological skin, likely will have an evolutionary advantage. The entire progress of media forms, from speech to writing, from tablets to papyrus, from digital memory in personal computers to the Internet can be considered the progress of this external memory, as detailed elsewhere [14]. What is equally important is not just the ability to retrieve a representation from a perfect (or at least "more perfect") memory, but the representation to be accessed by more than one individual at once.

Imagine the world to be inhabited by multiple individuals that can access the same representation. In almost all the original examples that Clark and Chalmers use in the Extended Mind Thesis, they deploy a single person sitting in front of a computer screen [7]. A more intuitive example would be two people using the Internet to both share a single representation. One could imagine Otto trying to find his way to the Museum of Modern Art, and instead of a notebook having a personal digital assistant with access to a map on the Web. Likewise Inga can having access to the exact same map via her personal digital assistant. Since both Otto and Inga are sharing the exact same representation and because they are both using it in the same manner, Inga and Otto can be said to share at least some of the same cognitive state, due to the fact that their individual cognitive states are causally dependent on accessing the same representation. This representation is the "same" precisely because of the perfect, digital memory of the computer. However, unlike the lone digital computer user of old, what the Web specializes in is allowing *everybody* to access the same set of representations.

The value of external representations comes with their accessibility, for an external representation that is not accessible when its needed cannot be used to enable online intelligence. It is precisely in order to solve this problem that Tim Berners-Lee proposed a World Wide Web as a universal information space [2]. The primary advantage of the Web is that every representation has a unique name, a URI[3] such as *http://www.example.org*. The Web allows each representation to be accessed when needed by using its unique name. This, combined with the fact that since the representations are digital and (at least can be) communicated in a lossless fashion, allows multiple simultaneous accessing of the exact same representation. Since the Web is a universal space of digital representations, two or more individuals can share the same representation simultaneously. Due to the Extended Mind hypothesis, two or more individuals can then, because of simultaneous access, share some of the same cognitive state.

6. THE WEB AS COLLECTIVE INTELLIGENCE

Much as computation has not remained static, neither has the Web. The Web, as originally conceived by its users, was just a collection of documents connected by hyperlinks, albeit one in a universal information space. This documents were mostly static, being authored and maintained by individuals. Although new pages and links could be added without resort to a centralized registry, the content of the Web was for the vast majority of users was not content that they actually created and added to in any meaningful manner. Within the last few years, a combination of easy-to-use interfaces for creating content and a large number of websites that prioritize the social and collaborative creation of content by ordinary users have taken off, leading to the phenomenon known as "Web 2.0," literally the next generation of the Web.[4]. This transition from the Web of static hy-

[2]In fact, many of us would need a map even without an impaired memory, which points to how widespread this phenomenon is.

[3]Originally the "Universal Resource Identifier," now a Uniform Resource Identifier as given in an updated specification [4]

[4]A term originally coined by Tim O'Reilly for a conference

perlinked web-pages to a more interactive and collaborative medium is more accurately described as a transition from a "Web of Documents" to a "Social Web" [12]. Paradigmatic examples of easy-to-use interfaces would be Google Maps (or even Google Earth), while a paradigmatic example of socially-generated content would be Wikipedia. Furthermore, increasingly these web sites are now being woven into the fabric of the everyday life of more and more people. How many people feel that their intelligence in increased when they have immediate access to a search engine to the Web, a massive encyclopedia available in a few seconds notice?

The Social Web then presents an interesting twist on the Extended Mind Hypothesis extension that we presented earlier. Again, Otto is using a web-page in his mobile phone to find his way to the Museum of Modern Art. While our previous example had Otto using the Web as ordinary Web users did years ago, simply downloading some directions and following them, we now add a twist. Imagine not only that Inga and Otto are using a map-producing Web site that allows users to add annotations and corrections, a sort of wiki of maps. Inga, noticing that the main entrance to the Museum of Modern Art is closed temporarily due to construction and so the entrance has moved over a block, adds this annotation to the map, correcting an error as regards where entrance of the Museum of Modern Art should be. This correction is propagated at speeds very close to real-time back to the central database behind the Web site. Otto is running a few minutes behind Inga, and due to this correction to the map being propagated to his map on his personal digital assistant, Otto can correctly navigate to the new entrance a block away. This (near) real-time updating of the representation was crucial for Otto's success. Given his memory issues, Otto would have otherwise walked right into the closed construction area around the old entrance to the Museum and been rather confused. This active manipulation with updating of an external representation lets Inga and Otto possess some form of dynamically-changing collective cognitive state. Furthermore, they can use their ability to update this shared external representation to influence each other for their greater collective success. In this manner, the external representation is clearly social, and the cognitive credit must be spread across not only multiple people, but the representation they use in common to successfully accomplish their behavior. Clark and Chalmers agree: "What about socially extended cognition? Could my mental states be partly constituted by the states of other thinkers? We see no reason why not, in principle" [7].

This leads us back full circle to the Web. For example, the collective editing of Wikipedia and its increasing use, allows its representations to be increasingly part of of the cognitive system of many people. As representations on the Social Web are updated by increasing numbers of people, each representation is increasingly brought into tighter coupling with both its target. As each representation is involved in this process of use and updating is brought into closer and closer cognitive updating with more and more individuals, the representations on the Web are brought into tighter and tighter coupling with what its users formerly considered their individual intelligence, and so leading to the phenomenon widely known as collective intelligence. Indeed, there are now problems as simple as navigating down the street, or organizing a social event, that many today would have difficult organizing without access to an interactive mapping Web service, or a social networking web site. As users contribute more and more content, the collective content of these web-pages becomes increasingly difficult to track down to individuals. Some of these Web-based tools for collective intelligence have no way to track down the original individual author, others like Wikipedia have sophisticated mechanisms in place to track individual contributions. However, as long as the contribution the collectively-built web page makes is the sum of more than an individual effort, then the credit must be placed upon the collective content, not the individual author. From the standpoint of the user of the representation, the credit must also not just be placed on the creator of the content, but the very technological infrastructure - ranging from the hardware of high-speed fibre optics and wireless routers to the software of protocol design and web server code - that enables the content of the collectively created web site to be delivered when it is needed. The credit for successfully creating and deploying the cognitive scaffolding is more collective than originally thought.

to describe the next generation of the Web

7. CONCLUSIONS

Having philosophers seriously move their research programmes into the nature of the Web will doubtless cause a paradigmatic shift in the debate over cognition and the Extended Mind, and thus more generally in the relationship between philosophy and the Web. A successful philosophy of the Web depends on taking an approach to the philosophical questions that remains grounded in the science and technology of the Web, including detailed rigorous inspection of empirical work we have not had the space to delve into here. However, it should be clear that a careful analysis of a wide interdisciplinary literature is necessary, a literature that extends beyond the traditional grounds of cognitive science and into studies of online communities, human-computer interaction, information retrieval, hypertext, and the Semantic Web. Although we have not answered all the questions that a philosophy of the Web should answer in order to provide answers to outstanding questions from the philosophy of mind and language, we have at least made a map of the territory for future research. Philosophy may be part of Web Science after all.

References

F. Adams and K. Aizawa. The bounds of cognition. *Philosophical Psychology*, 14:43–64, 2001.

T. Berners-Lee. IETF RFC 1630 Universal Resource Identifier (URI), 1994. http://www.ietf.org/rfc/rfc1630.txt (Last accessed on May 3rd 2008).

T. Berners-Lee. Message on www-tag@w3.org list, 2003. http://lists.w3.org/Archives/Public/www-tag/2003Jul/0158.html (Last accessed on May 20th 2008).

T. Berners-Lee, R. Fielding, and L. Masinter. IETF RFC 3986 Uniform Resource Identifier (URI): Generic Syntax, January 2005. http://www.ietf.org/rfc/rfc3986.txt (Last accessed on April 2th 2008).

T. Berners-Lee, J. Hendler, and O. Lassila. The Semantic Web. *Scientific American*, 284(5):35–43, May 2001.

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

A. Clark and D. Chalmers. The Extended Mind. *Analysis*, 58(1):7–19, 1998.

H. Dreyfus. *What Computers Still Can't Do: A critique of artificial reason*. MIT Press, Cambridge, Massachusetts, USA, 1979.

H. Dreyfus. *On the internet*. Routledge, New York, 2001.

L. Floridi. Web 2.0 vs. the semantic web: A philosophical assessment. *Episteme*, 6(1):25–37, 2009.

H. Halpin. Philosophical Engineering: Towards a Philosophy of the Web. *APA Newsletter on Philosophy and Computers*, 7(2):5–11, 2008.

P. Hoshka. CSCW research at GMD-FIT: from basic groupware to the social web. *ACM SIGGROUP Bulletin*, 19(2):5–9, 1998.

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

R. Logan. *The Sixth Language*. Stoddart Publishing, Toronto, 2000.

J. McCarthy and P. Hayes. Some philosophical problems from the standpoint of Artificial Intelligence. In B. Meltzer and D. Michie, editors, *Machine Intelligence*, volume 4, pages 463–502. Edinburgh University Press, 1969.

K. O'Hara and W. Hall. Trust on the web: Some web science research challenges. *UoC Papers: E-Journal on the Knowledge Society*, (7), October 2008.

P. R. Smart, P. C. Engelbrecht, D. Braines, M. Strub, and C. Giammanco. The network-extended mind. In D. Verma, editor, *Network Science for Military Coalition Operations: Information Extraction and Interaction*. IGI Global, Hershey, Pennsylvania, USA, 2010.

F. J. Varela, E. Thompson, and E. Rosch. *The Embodied Mind: Cognitive Science and Human Experience*. MIT Press, Cambridge, Masschusetts, USA, 1993.

M. Wheeler. *Reconstructing the Cognitive World: The Next Step*. MIT Press, Cambridge, Massachusetts, USA, 2005.

M. Wheeler. The Fourth Way: A comment on Halpin's 'Philosophical Engineering'. *APA Newsletter on Philosophy and Computers*, 8(1):9–12, 2008.

Y. Wilks. A personal memoir: Margaret Masterman (1910-1986). In M. Masterman, editor, *Language, Cohesion, and Form*. Cambridge University Press, Cambridge, United Kingdom, 2005.