

Philosophical Engineering: Towards a Philosophy of the Web

Harry Halpin
H.Halpin@ed.ac.uk
University of Edinburgh

Abstract:

The Web is commonly considered the most significant computational phenomenon to date. However, the Web itself has received scant attention from philosophy, being best regarded as a mere engineering artifact. Furthermore, the efforts to evolve the Web into the Semantic Web are viewed with suspicion by most philosophers as a return to Cartesian artificial intelligence. I argue that these widely held viewpoints are incorrect, and that the Web succeeds because of its design principles that distinguish it from both previous hypertext systems and knowledge representation systems in classical AI. Furthermore, the Web embodies the logical conclusion of Clark's Extended Mind thesis since it allows multiple individuals to access and manipulate the same representation, so offering the ultimate in cognitive scaffolding. This undermines the notion of individual intelligence at the heart of Cartesian artificial intelligence and presents a challenge to the role of representations as given in the recent wave of neo-Heideggerian focus on embodiment. Taking the Web seriously moves the primary focus of philosophy away from the role, or lack thereof, of internal representations to external representations. The Web is then properly understood as the creation and evolution of external representations in a universal information space. Berners-Lee calls this "philosophical engineering," and it has surprising connections to neo-Fregeanism, anti-realism, and other long-standing philosophical debates.

The Web as Philosophy

The reigning model of computation usually considered by philosophers debating artificial intelligence is the lone Turing Machine manipulating representations in an existential void. The most controversial proposition of computer science that captures the imagination of philosophy is whether or not this model can be authentically intelligent. Under heavy attack from philosophers ranging from Dreyfus to Wheeler, this hypothesis of classical artificial intelligence is showing wear and tear, if not total defeat. From the ashes of what has been termed "Good Old Fashioned Artificial Intelligence" has arisen a neo-Heideggerian theory of embodiment as based on work in dynamical systems, robotics, and artificial life, inspired by the audacious claim of Brooks that intelligence does not require representations (Brooks, 1991). While impressive, the vast majority of work in computing is not the building of embodied robots capable of simulating emotions. While the sheer ambition of embodied artificial intelligence grabs the imagination of philosophers, a revolution of everyday life has emerged due to the advent of the Web. The Web, and the complex web of human and computer interaction it engenders, is already supplanting the isolated digital computer in everyday life, and so will eventually supplant it as the basis for philosophical analysis of the nature of computation and the mind. We argue that the Web represents a fundamental turning point not just in computational systems but in the philosophy of information and the mind. The Web signals a return of representationalism of a distinctly different kind than that

formerly theorized about by artificial intelligence, “embodied” or otherwise.

In his “One Billion Lines of C++,” Smith notes that the models of computing used in philosophical debates over intelligence, representation, embodiment, and consciousness ignore the majority of existing computers by framing the debate as if it were between logic-based symbolic reasoners and some alternative ranging from neural networks to epigenetic robotics (1997). As Smith points out: “it is impossible to make an exact estimate, but there are probably something on the order of 10, or one hundred billion lines of C++ in the world. And we are barely started. In sum: symbolic AI systems constitute approximately 0.01% of written software” (Smith, 1997). The same small fraction likely holds true of “non-symbolic AI” computational systems such as robots, artificial life, and connectionist networks. Numbers by themselves hold little intellectual weight, for one could always argue that the vast majority of computational systems are simply philosophically uninteresting. However, trends in philosophy contradict this intuition. The turn in philosophy away from artificial intelligence and linguistics to what has been termed a “philosophy of information” demonstrates that this wider class of computational systems are finally having an impact (Floridi, 2004).

The most significant computational system to date is the World Wide Web, described by its inventor Tim Berners-Lee as “a universal information space” (Berners-Lee, 1998). Due to its hegemonic role today, understanding the principles that distinguish the Web from other computational systems should be a goal of the philosophy of the information. To articulate what I term the “philosophy of the Web,” one needs to first clarify whether or not the Web is continuous or in conflict with the development of current trends in philosophy, cognitive science, and artificial intelligence. This is not an easy endeavor at first glance, for the Web is nothing if not a robustly representational system, and a large amount of research on the Web focuses on how to enable increasingly powerful and flexible forms of representations. One development Berners-Lee calls the “Semantic Web” consists of standardizing the logical encoding of representations so that they may be directly usable by machines. There is suspicion among many of the now successful rebels against classical artificial intelligence that “Good Old Fashioned” artificial intelligence has not disappeared but instead has been born anew on the Semantic Web.

Beyond Neo-Heiddegerian Embodiment

It is difficult to summarize the hypothesis made by proponents of embodied artificial intelligence, yet this must be done in order to understand if this hypothesis is in fact continuous with the Web. In contrast, in defining classical artificial intelligence, Brian Cantwell Smith stated the Knowledge Representation Hypothesis: “Any mechanically embodied intelligent process will be comprised of structural ingredients that a) we as external observers naturally take to represent a propositional account of the knowledge that the overall process exhibits, and b) independent of such external semantical attribution, play a formal but causal and essential role in engendering the behavior that manifests that knowledge” (Smith, 1985). One compelling story put forward by the philosophers of embodied artificial intelligence is by Michael Wheeler (Wheeler, 2005). Instead of phrasing a negative critique of classical artificial intelligence like Hubert Dreyfus, Wheeler turns his Heideggerian analysis into a positive programme for

embodied artificial intelligence. In order to contrast classical artificial intelligence with its embodied alternative, Wheeler produces three Cartesian claims that he believes underlie classical artificial intelligence (Wheeler, 2005):

1. The subject-object dichotomy is a primary characteristic of the cognizers ordinary epistemic situation.
2. Mind, cognition, and intelligence are to be explained in terms of representational states and the ways in which such states are manipulated and transformed.
3. The bulk of intelligent human action is the outcome of general purpose reasoning processes that work by retrieving just those mental representations that are relevant to the present behavioral context and manipulating and transforming those representations in appropriate ways as to determine what to do.

Wheeler then states that “word on the cognitive-scientific street is that classical systems have, by and large, failed to capture in anything like a compelling way, specific styles of thinking at which most humans naturally excel” (2005). Wheeler then contrasts the Cartesian assumptions of classical artificial intelligence with Heidegger in order to formulate the philosophical principles of embodied artificial intelligence (Wheeler, 2005):

1. **The primacy of online intelligence:** The primary expression of biological intelligence, even in humans, consists not in doing math or logic, but in the capacity to exhibit...online intelligence...a suite of fluid and flexible real-time adaptive responses to incoming sensory stimuli.
2. **Online intelligence is generated through complex causal interactions in an extended brain-body-environment system:** Online intelligent action is grounded not in the activity of neural states and processes alone, but rather in the complex causal interactions involving not only neural factors, and also additional factors located in the non-neural body and the environment.
3. **An increased level of biological sensitivity:** Humans and animals are biological systems – *and that matters for cognitive science*.
4. **A dynamical systems perspective:** Cognitive processing is fundamentally a matter of state space evolution in certain kinds of dynamical systems.

These assertions we call the *neo-Heideggerian framework* of embodied cognitive science. Wheeler and other philosophers then return to the old question of whether or not there is room for internal representations in cognitive science. Wheeler argues that there is a limited role for internal representations to play, and work in embodied cognitive science should no longer only tout that the world is its own best model (Brooks, 1991). Even work in robotics has shown that internal representations are incredibly important when “the world is not enough” and that training via simulation can be effective replacement for an often dangerous world (Grush, 2003). The world may be its own best model, but it can get you killed. Therefore, some limited form of representations may

have an evolutionary advantage for those that can use them, and may even fit within the neo-Heideggerian framework (Grush, 2003). This is a small victory at best for representations, for representations are cast from their long-standing spotlight in classical artificial intelligence into a secondary role in embodied cognitive science.

What are Representations?

The very idea of representation is usually left under-defined as a vague “standing-in” intuition. The classic definition of a symbol from the Physical Symbol Systems Hypothesis is the genesis of this intuition regarding representations (Newell, 1980):

“An entity X designates an entity Y relative to a process P , if, when P takes X as input, its behavior depends on Y .”

There are two keys to this definition. First, the concept of a representation is grounded in the behavior of a process. Thus, what precisely counts as a representation is never context-free. Second, the representation simulates action at a distance on what is being represented: “This is the symbolic aspect, that having X (the symbol) is tantamount to having Y (the thing designated) for the purposes of process P ” (Newell, 1980). This definition seems to have an obvious point of conflict with the neo-Heideggerian agenda, for it reflects a “subject-object dichotomy” due to its presupposition of three ontologically distinct entities. To return to the Physical Symbol Systems Hypothesis's definition, let us call P the “subject” that is using the representation. Let us call X the “representation” and Y the “object.” Therefore, the subject-object dichotomy is actually present three times in the Physical Symbol Systems Hypothesis: the subject, the object, and the representation. The dichotomy is present between the subject and its object, the object and its representation, and even perhaps the subject and its representation. To the extent that these distinctions are ontologically held a priori, then the definition is hopelessly Cartesian.

The only way to escape this Cartesian trap is to give a description of how representations arise without the a priori subject-object dichotomy. This is precisely what Brian Cantwell Smith proposes through his process of *registration* (1996). Smith starts with the classic example of a frog tracking a gadfly across the sky. The frog sees the gadfly, and begins tracking the gadfly with its eyes as it flies. The frog and the gadfly are both physically connected via light-rays. Borrowing an analogy from physics, Smith notes that everything is composed of non-distinct fields of energy, so it would be ontologically wrong to talk about a frog, a gadfly and light as individual objects. All that exists are fields that are physically connecting and disconnecting from each other to greater or lesser degrees. At the moment of tracking, one can speak of the frog and gadfly as a single connected field. When the gadfly goes behind a tree, and then when the fly emerges from the other side of the tree, the eyes of the frog are not focused on the point the gadfly was at before it went behind the tree, but the point the gadfly would be at if it continued on the same path. Although this is a simple trick capable of being accomplished without any full-blooded internal representation, Smith believes this simple case builds the foundation for the eventual emergence of representation. In the language of fields, the fields separate into an *o-region* disconnected from the *s-region*. The s-region is distinguished from the o-region by virtue of the s-region's attempt to “track” the o-

region, its attempt to remain connected in “a long-distance coupling against all the laws of physics” (Smith, 1996). The s-region eventually stabilizes as the subject and the o-region as an object, with considerable work on at least the subject's side. This work manifests itself as the creation of a representation that the subject maintains of the object. The subject, the representation of the object, and the object itself are the final result of registration. Furthermore, the subject and object are not a priori distinct, but co-constitute each other. The distinction between subject and object is given by the use of a representation of the object by the subject.

In order to explicate what precisely the subject must possess in order to track the object via a representation, we rely on Rocha and Hordijk's notion of *dynamically incoherent memory* (2005). “Dynamically incoherent” means the memory is not changed by any dynamic process it initiates or encounters. The term “dynamically incoherent” is misleading, for most people would say “dynamically incoherent” actually means the maintaining of *coherence* against “the vagaries and vicissitudes, the noise and drift, of earthy existence,” as Haugeland would say (1981). In other words, somehow the object must have a memory that allows it to store a representation for some period of time with a degree of fidelity and reliability. Of course, this is precisely what *digital* memory is good for.

However fuzzy the details of Smith's story about representations may be, what is clear is that instead of positing the subject, object, and representation a priori, they are introduced as products of a temporal process. This process is at least theoretically “non-spooky” since the entire process is capable of being grounded out in physics without any “spooky” action at a distance. To be grounded out in physics, all changes must be given in terms of contact in space-time, or in other words, “local” contact. In this way, representations are “a way of exploiting local freedom or slop in order to establish coordination with what is beyond effective reach” (Smith, 1996). In order to clarify Smith's registration and improve the definition of the Physical Symbol Systems Hypothesis, we consider this entire story to be a *representational cycle* (Halpin, 2006):

1. **Presentation:** Process S is in local contact with process O . S is the s-region, that evolves into the subject, while O is the o-region that evolves into the object.
2. **Input:** The process S is in local contact with coherent memory R . An input procedure of S puts R in correspondence with some portion of process O . This is entirely non-spooky since S and O are in local contact with R . R evolves into the representation.
3. **Separation:** Processes O and S change in such a way that the processes are non-local.
4. **Output:** Due to some local change in process S , S uses its local contact with R to initiate local behavior that depends on R for success.

So we have constructed an ability to talk about representations while not presupposing that intelligent behavior depends on internal representations or that representations exist a priori at all. Representations are only needed when the relevant intelligent behavior requires some sort of distal co-ordination. In this manner, representations, if not representationalism, is continuous with the neo-Heideggerian

agenda. Representations are not a Cartesian metaphysical assumption, but arise over time in a way that is not only coherent with physics but with the neo-Heideggerian programme. Representations exist as part of a rich temporal dynamic that does not presuppose a Cartesian subject-object dichotomy, instead being based on contingent and temporary object-subject dichotomies.

From the Extended Mind to External Representationalism

The success of Smith's argument lies also in its ability to phrase the creation of representations that are neither internal or external to a particular subject. Furthermore, just because the coherent memory that serves as a representation is at some point in local contact with a subject, it does not mean that it must always be tied to the subject. In other words, the representation does not have to be “in the head” of the subject. In this case, we can divide the world of representations into two types, those that are internal to the subject and those that are external to the subject. Although this argument undermines the Cartesian use of “internal” and “external” in general, for the rest of this argument we will just use the term “internal representation” to designate representations implemented biologically inside a human subject, as traditionally defined by the bounds of the skin. We will use “external representation” to designate representations that are implemented outside what traditionally considered outside the biological body of a human subject.

One of the tenets of the neo-Heideggerian programme put forward by Wheeler is that “online intelligence is generated through complex causal interaction in an extended brain-body-environment system.” (2005). There is no reason why representations can not be part of that environment, since in lieu of Smith's story of registration we have rephrased representations as not necessarily being internal. In fact, we can remain agnostic as regards to the possibility of whether or not internal representations are necessary or even used by a human subject and leave the debate over internal representations is a purely empirical question best left to cognitive science. These representations in the environment can even be crucial to intelligence, leading to what Clark and Chalmers call “an active externalism, based on the active role of the environment in driving cognitive processes,” their Extended Mind Thesis (1999).

To explain their Extended Mind thesis, Clark and Chalmers present introduce us to the charming Otto, a man with an impaired memory who navigates about his life via the use of his notebook (1999). Otto is trying to navigate to the Metropolitan Museum of Modern Art in New York City from his house. In order to successfully arrive at the museum, Otto needs a map which is in some correspondence with the world he must navigate through in order to successfully get to the museum. We can imagine Otto has a map in his notebook to the Museum of Modern Art made for the precise purpose of navigating individuals to the museum. Otto can get to the museum with the map, and without it he would be lost. It is hard to deny that a map is representational in the sense we have presented above. In this regard, external representations do exist in the environment of an agent and can drive the cognitive processes of an agent in a similar fashion to the way that classical artificial intelligence assumed internal representations did. Interestingly enough, 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” (1999). In this regard, the Extended Mind thesis undermines the strict division between internal and external of the subject and its representation in a way that is compatible with the neo-Heideggerian philosophical framework.

The Extended Mind in combination with the representational cycle explains the widespread success of digital computing. The analog memories of humans are notoriously bad at maintaining any sort of coherence in front of the push and pull of everyday life. This means that humans who discover a way to maintain representations outside of their biological skin may prove to have an evolutionary advantage over those who don't have such an ability. The maintenance of coherent memory requires considerable work. In digital memory, this is exemplified in how voltages must be carefully regulated to maintain the coherence of computer memory. With the advent of digital memory, where representations can be safely encoded and recoded at a level of abstraction that is resistant to change, a new medium for representations has been created that is in general more suited for representations than biological memory. This may explain the strange fascination of classical artificial intelligence for digital computing and internal representationalism, even though many representation-heavy tasks studied by classical artificial intelligence constitute “psychological arenas in which most humans perform rather badly, and in which most other animals don't perform at all” (Wheeler, 2005). We might want to amend this observation. Classical artificial intelligence specialized in tasks that humans perform rather badly at when not aided by machines with digital memory. With the advent of digital memory and computation humans can now be successful at all sorts of things that humans and other animals without the cognitive scaffolding of digital technology would fail at normally.

To press upon the Extended Mind thesis, imagine the world to be inhabited by multiple subjects that can access the same representation. In almost all the original examples that Clark and Chalmers use in the Extended Mind Thesis, they use a lone person sitting in front of a computer screen. This ignores the use of multiple people using the computer as a collaborative and communications tool. The obvious 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 mobile telephone with access to a web page that has a map. One could also imagine Inga, who may not have Alzheimer's but nonetheless cannot remember her way to the Museum, having access to the same map via her personal digital assistant's Internet access. Since they are sharing the representation and their behavior is normatively successful based on its use, Inga and Otto can be said to partially share the same cognitive state.

To push the Extended Mind thesis even further, imagine not only that Otto and Inga are using a web page with a map, but a web page that allows users to add annotations to the map. The web page is updated with the annotations for all to use in near real time. 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 should be. Luckily, with this change to the map Otto can now find the entrance to the museum,

while without it he would have been hopelessly lost. This active manipulation of a representation lets Inga and Otto partially share a dynamic cognitive state and collaborate for their greater collective success. Their shared cognitive process is functioning not via telepathy but via shared external representations that are universally accessible over the Web. Clark and Chalmers agree that cognition can be socially extended: “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” (1999). This socially extended cognition can be accomplished via shared external representations. In fact, it is precisely to these ends that the constant work on the Web is heading, especially the more interactive and collaborative Web technologies.

The Principles of the Web

Unlike classical artificial intelligence's emphasis on internal representations, the Web is primarily concerned with external representations. In contrast to the champions of the embodied mind, the Web is concerned with exactly what digital, external, representations can be used to maximize intelligence. It is perhaps an irony that the embodied mind is making such gains while human society and intelligence seem increasingly disembodied. The value of external representations comes with their accessibility, for an external representation that is unaccessible can not enable online intelligence. It is precisely in order to solve this problem of the accessibility of external representations that Tim Berners-Lee proposed the Web as “a universal information space” (1998). The primary advantage of the Web is that representations can be accessed via a URI¹ so that it may link with any other representation. Furthermore, a URI identifies a *resource*, which may in turn host multiple representations of the same object. From this foundation in URIs, a number of standardized protocols for accessing representations² and formats for encoding representations³ provide the foundation of the Web. From the perspective of philosophy, it is interesting that traditional philosophy has in general taken inspiration from artificial intelligence, although it is only one branch of research in computing. Not as well-known as the lineage of artificial intelligence, the Web has an historical lineage in Douglas Engelbart's “Human Augmentation Framework” (1962), a parallel and some times competing programme to artificial intelligence. The Web inherits much inspiration and even design from the work of Engelbart's NLS (oNLine System). More and more recently his system and ideas have began to receive attention from philosophers (Cooksey, 2006).

What would be philosophically important would be if certain kinds of external representations were found to be amendable to facilitating intelligence. To answer this question, the exact characteristics of the external representations used by the Web need to be articulated. While their articulation requires far more detailed explanation, the principles of the Web may be defined as:

- 1. The Principle of Universality:** Everything can be given a URI, and therefore everything can be identified on the Web in order to retrieve a representation of it.
- 2. The Principle of Inconsistency:** A representation can be linked to URIs, and these links can be inconsistent (i.e. resulting in a URI that does not exist, returning something like the infamous “HTTP 404 Not Found” error).

3. The Principle Self-Description: All representations should be linked to URIs that describe the meaning of the URI.

4. The Principle of Least Power: A representation should be described in the least powerful but adequate format.

5. The Principle of the Open World: The number of URIs, and so representations, can always increase.

These principles are a systemization that takes inspiration from the work of Berners-Lee and other designers of the Web (1998). Taken together, these principles are all independent. One can describe a possible web in which some combination of these principles does not hold. However, at this time we hold the conjunction of these principles minimally describe the philosophically interesting features of representations on the Web, and are so a good starting point to investigate the philosophical basis of the Web. Upon the surface, it does not appear that these principles may have philosophical weight, yet appearances can be deceiving. The connections between these principles and long-standing philosophical debates are productive and deep.

The first principle of universality posits that literally everything can be given a unique identifier, and as such relates to work on how identifiers like names are established. This is related to the debate between Kripke's causal theory of names and neo-Fregean descriptivist accounts of names. It seems that the “meaning” of a web page on the Web can be either given by the person who controls its URI, which is established in a clear “baptizing” process through domain name registration, which is philosophically similar to Kripke's theory of proper names (Luntley, 1999). Alternatively, one could imagine that the “meaning” of a web page is given by the representations it returns in a neo-Fregean manner, so that even if the owner of the URI thought it was about blue cheese, if the URI returned pictures of the moon people would be correct in assuming that the particular URI was about a moon. Other principles point to areas that are under-explored by philosophy. The Principle of Self-Description demands a coherent story about what “self-description” even means. Self-description may be related to notions explored by Kolmogorov's algorithmic information theory. The Principle of Least Power seems to demand a notion of power not easily mapped to traditional notions of power like Turing-completeness and the Chomsky Hierarchy.

The clearest correspondence to classical problems in artificial intelligence that external representations on the Web provoke is to problems of reasoning, but even here there are crucial differences like the Principle of Inconsistency. Unlike previous hypertext systems, on the Web any representation can be linked to any other representation without a centralized database of links. In philosophical terms, if a link represents a logical predicate and a web-page is a fact, then there is no principle of consistency on the Web. Long considered problematic for logic and philosophy, inconsistency is elevated to the status of a defining principle on the Web. Tolerance for inconsistency is precisely what removes the largest limiting factor of previous hypertext systems, since it allows users of the Web to link representations to URIs in whatever fashion they find most useful without asking permission. Combined with universality, this principle furthers the “network

effect” where the value of each representation grows in proportion to the size of the Web, since any representation may be linked to any other representation on the Web. In a historical note, it should be remembered that the original academic paper in which Berners-Lee attempted to present the idea of the Web was rejected precisely because of this inconsistency, yet it is inconsistency that allowed the Web to grow at such an astounding rate. Search engines like Google create post-hoc a centralized index of the Web, yet they are always behind the perpetual growth of the Web. For at every moment Google crawls a web-page, another web page can appear.

The Principle of the Open World states that the number of representations is always increasing. Therefore, unlike classical mathematics, it is difficult to ever say that a fact is false on the Web as a whole, since the set of representations one is reasoning about is ever-increasing. In order to pursue a strategy that says any particular fact is false or even unknown, one must somehow draw a closed boundary over the Web, which violates the Web being an open system. Furthermore, since inconsistency is allowed, in an open system like the Web where diverse agents are always adding new “facts” (new web-pages being created, links being added, new logical assertions made), it is untenable even with an arbitrary closed portion of the Web to say a fact is false, for it may merely be inconsistent. This leads to another violation of traditional reasoning in classical artificial intelligence that tried to mimic human-reasoning by developing non-monotonic inference, because on the ever-increasing Web where the results of an inference are just another part of the Web, all inference should be monotonic. Due to this lack of a classical notion of true and false, the notion of truth on the Web can only be saved through intuitionism, in which truth is given by proof and any proof must take the form of a constructivist proof that does not rely on the Law of the Excluded Middle. So nothing if strictly false, “truth” is only what can be proven from a given set of facts selected from the Web and “false” that which cannot be proved from those particular set of facts. This exemplifies the insight of anti-realism of Dummett as applied to external representations: the debate between intuitionism and Platonism in mathematics has much wider philosophical repercussions, for on the Web it has engineering repercussions (1959). It leads to an abandonment of the closed world solution to the Frame Problem in favor of a proof-theoretic notion of truth that can survive the open world of the Web.

Conclusions: The Semantic Web Reconsidered

By virtue of being an open system of external representations, the Web has a number of principles that directly conflict with classical artificial intelligence, yet due to its heavy use of representations, the Web has naturally as it evolved attracted a strange affinity with classical artificial intelligence. Search engines like Google occupy the space in popular imagination that the all-knowing robotic brains of classical artificial intelligence once did. The “Semantic Web” of Berners-Lee is at the heart of this fear of the return of AI. At the very first World Wide Web Conference, Berners-Lee announced plans to move the Web away from mere hypertext to a web of meaning: “To a computer, then, the web is a flat, boring world devoid of meaning...this is a pity, as in fact documents on the web describe real objects and imaginary concepts, and give particular relationships between them...adding semantics to the web involves two things: allowing documents which have information in machine-readable forms, and allowing links to be created with relationship values.”⁴

Long thought to be vanquished by the success of the neo-Heideggerian embodied mind, the spectre of classical artificial intelligence has seemingly returned on the Web in the form of the Semantic Web. The first step in this effort was the creation of a knowledge representation language for the Web. At the beginning the only research community involved heavily was the classical artificial intelligence community. Prominent champions of classical artificial intelligence such as Pat Hayes helped create its formal semantics (Hayes, 2004). Yet the Semantic Web is not classical artificial intelligence, it is in fact something new, as our principles above demonstrated. Classical solutions often do not work on the Web due to scalability issues and the open-ended nature of the Web. Due to these factors, formerly obscure areas of research like description logic, which guarantee decidability over open-ended data sets, are coming to the forefront of research in the Semantic Web. The revival of classical artificial intelligence on the Web makes perfect sense, since in the world of carefully engineered external representations, skills that humans lack but computers have in spades such as “logic-based reasoning or problem-solving in highly structured search spaces” can be crucial (Wheeler, 2005). The type of problem the Semantic Web is meant to deal with is structured data-sharing and inference, which are more mundane than creating intelligence but perhaps just as useful in aiding human intelligence. Furthermore, with large reams of data, statistical methods often originating in artificial intelligence have even more so than logic-based artificial intelligence proven to be crucial to the success of the Web, as search engines like Google show. Biological sensitivity makes little sense in the world of representations, for the question is not what can an intelligent human do, but how can computers complement an intelligent human. In the words of Andy Clark, we are “human-technology symbionts: thinking and reasoning systems whose minds and selves are spread across biological brain and non-biological circuitry.” This symbiosis is done increasingly in practice today by the universalizing power of the Web.

Representations on the Web are machine-encoded, external, radically public, and reliant on a complex representational network of links. The meaning of these representations are ultimately grounded in their use by machines and humans. This is simply different from previous views of representations as primarily human-encoded, internal, private, governed by their truth conditions, and ultimately grounded in the neural activity of humans. While we would not dispute the possibility of the existence of internal representations, traditional artificial intelligence and analytic philosophy may be overemphasizing the role of internal representations in intelligence and underestimating the value of external representations. The wildfire growth of the external representations as fueled by the Web should give us doubt about any anti-representationalist arguments about human intelligence. In fact, what the Web seems more concerned with is reducing what Floridi calls *ontological friction*: “the forces that oppose the flow of information within (a region of) the infosphere and, hence, (as a coefficient) to the amount of work and effort required to generate, obtain, process, and transmit information in a given environment” (2007). It is precisely this minimization of ontological friction that the Web, and the Semantic Web, is trying to do. What is even more intriguing is the notion that at this point in time, non-biological representations may be more and more in the evolutionary driving seat. If this is indeed the case, then we have good reason to believe that the design and engineering of these external representations is no trivial task. Tim Berners-Lee has argued that “we are not analyzing a world, we are building it. We are not

experimental philosophers, we are philosophical engineers.”⁵ One insight of the philosophy of the Web is that these representations on Web do have philosophical significance, and philosophers are needed to help clarify their foundations. Another insight is that on the Web representations can indeed change the world.

Works Cited

- Berners-Lee, T. “Axioms of Web Architecture” 1998, <<http://www.w3.org/DesignIssues/Principles.html>> (22 Dec 2007).
- Brooks, R. “Intelligence without Representation.” *Artificial Intelligence* 47 (1991): 139-159.
- Clark, A. and D. Chalmers. “The Extended Mind,” *Analysis* 58 (1999): 10-23.
- Clark, A., 2003. *Natural-Born Cyborgs: Minds, Technologies and the Future of Human Intelligence*. Oxford, UK: Oxford University Press.
- Cooksey, M. “Exploring the Engelbart Hypothesis: A Philosophical Investigation.” *APA Newsletter on Philosophy and Computers*. 5 (2006).
- Dummett, M. “Truth.” *Proceedings of the Aristotelian Society*, 59 (1959): 141-162.
- Engelbart, D. *Augmenting Human Intellect: A Conceptual Framework*. Stanford Research Institute, Menlo Park, Ca., October 1962.
- Floridi, L., ed., 2004. *Philosophy of Computing and Information*. Malden, MA: Blackwell Publishing.
- Floridi, L. “Understanding Information Ethics,” *APA Newsletter on Philosophy and Computers*. 7 (2007).
- Grush, R. “In defence of some 'Cartesian' assumptions concerning the brain and its operation.” *Biology and Philosophy* 18 (2003): 53-93.
- Halpin, H., 2006. “Representationalism: The Hard Problem for Artificial Life.” In *Tenth Conference on Artificial Life*. Bloomington: USA: Indiana University.
- Haugeland, J. “Analog and Analog,” *Philosophical Topics* 12 (1981): 213-226.
- Hayes, P. “RDF Semantics” 2004, W3C Recommendation <<http://www.w3.org/TR/rdf-mt/>> (Dec 18th 2007).
- Luntley, M., 1999. *Contemporary Philosophy of Thought*. Oxford, UK: Blackwell.
- Newell, A. “Physical Symbol Systems,” *Cognitive Science*, 4 (1980): 135-183.
- Rocha, L. and Hordijk, W. “Material Representations: From the Genetic Code to the

Evolution of Cellular Automata.” *Artificial Life*, (2005) 11: 189-214

Smith, B.C., 1985. “Prologue to reflection and semantics in a procedural language.” In *Readings in Knowledge Representation*, edited by R.J. Brachman and H.J. Levesque. 31-40. Morgan Kaufmann, 1985.

Smith, B.C., 1996. *On the Origin of Objects*. Cambridge, MA: MIT Press.

Smith, B.C., “One Billion Lines of C++.” *LeHigh CogSci News* 10 (1997)
<<http://www.lehigh.edu/~incog/v10n1/smith.htm>> (Dec. 18th 2007).

Wheeler, M., 2005. *Reconstructing the cognitive world: the next step*. Cambridge, MA: MIT Press.

¹ Originally the ``Universal Resource Identifier," now a Uniform Resource Identifier, such as *<http://www.example.org>*.

² The most prominent being HTTP (HyperText Transfer Protocol), although FTP (File Transfer Protocol) is also well-known.

³ Such as HTML (Hypertext Markup Language) and XML (Extensible Markup Language)

⁴ See <http://www.w3.org/Talks/WWW94Tim/> for the complete slides from Berners-Lee's WWW 1994 announcement.

⁵ See <http://www.xml.com/pub/a/2003/07/23/deviant.html> for a synopsis of the argument between Hayes and Berners-Lee, as well as a wider critique. .