

CHAPTER 1

TOWARD A PHILOSOPHY OF THE WEB: FOUNDATIONS AND OPEN PROBLEMS

ALEXANDRE MONNIN AND HARRY HALPIN

Introduction

What is the philosophical foundation of the World Wide Web? Is it an open and distributed hypermedia system? Universal information space? How does the Web differ from the Internet? While the larger ecology of the Web has known many a revolution, its underlying architecture in contrast remains fairly stable. URIs (Uniform Resource Identifiers), protocols like HTTP (HyperText Transfer Protocol), and languages such as HTML (HyperText Markup Language) have constituted the carefully evolved building blocks of the Web for more than two decades. As the particular kind of computing embodied by the Web has displaced traditional proprietary client-side applications, the foundations of Web architecture and its relationship to wider computing needs to be clarified in order to determine the Web's roots and boundaries, as well as the historical reasons for its success and future developments. Crafting a philosophy of the Web is especially urgent, as debate is now opening over the relationship of the Web to platform computing on mobile devices and cloud computing.

The scope of the questions that the philosophy of the Web provokes is quite wide-ranging. These questions begin with the larger metaphysical issue of whether or not there are unifying principles underlying the architecture of the Web that justify the existence of a philosophy of the Web. Tim Berners-Lee, widely acclaimed as the inventor of the Web, has developed in his design notes various informal reflections over the central role of URIs (Uniform Resource *Identifiers*, previously *Locators*) as a universal naming system, a central topic in philosophy since at least the pioneering works of Barcan Marcus. URIs such as <http://www.example.org/> identify anything on the Web, so the Web itself can be considered the space of all URIs. Thus, in brief, we would say that there is indeed at least one unifying principle to the architecture of the Web, that of URIs. The

various architects of the Web, including Berners-Lee, made a number of critical design choices, such as creating a protocol-independent universal naming scheme in the form of URIs as well as other less well-known decisions, such as allowing links to URIs to not resolve (leading to the infamous “404 Not Found Message,” a feature not allowed in previous hypertext systems) that—little to the knowledge of everyday users of the Web—do form a coherent system, albeit one that has not yet been explicated through a distinctively philosophical lens.

A critic could easily respond that there is no *a priori* reason any particular technology deserves its own philosophy. After all, there is no philosophy of automobiles or thermostats. Why would one privilege a philosophy of the Web over a philosophy of the Internet? These questions can be answered by looking at the nature of the design choices made in the formation of the Web: namely, in so far as the Web is based on URIs, the architecture of the Web exists on the level of naming and meaning, both of which are central to semantics and so are traditionally within the purview of philosophy. What the Web adds to the traditional philosophical study of natural language is both the *technically* engineered feat of a universalizing naming scheme in the form of URIs and the fact that such names can be accessed to return concrete bits and bytes, a distinctive feature of naming on the Web. However, the Web itself is agnostic over how the concrete low-level bits that compose something like a web page are transmitted across the network in response to an access request to a URI, as this is determined by protocols such as the Internet’s TCP/IP (Transmission Control Protocol/Internet Protocol). Thus, the Web can be considered an abstract information space of names above the networking protocol layer, up to the point that it could have been (or could still be) built on top of another networking protocol layer (such as OSI [Open Systems Interconnection] or the “Future Internet”). Likewise, the Internet can also host applications other than the Web that do not use URIs, such as peer-to-peer file sharing or the Web’s early rivals (the Gopher system, for instance). So in response to our critic, the Web does have its own architecture, and—unlike the case with automobiles and even the Internet—this architecture uncontroversially deals with philosophical concepts of naming and meaning, and this justifies the existence of a philosophy of the Web, at least insofar as names and meaning on the Web differ from natural language (or the philosophical way to conceptualize it!), a topic worthy of further exploration (Monnin 2012a).

The Web is not all protocols and naming schemes; it is also a wide-ranging transformation of our relationship to the wider world “out there,” to the ontology of the world itself. It is precisely this engineering aspect that makes the philosophy of the Web differ qualitatively from traditional philosophy of language, where it has been assumed that natural language is (at least for philosophical purposes) stable and hence “natural.” In

contrast, the nature of the growth of both the Web and digital technologies undoubtedly calls into question the contemporary transformation of our entire form of life. Bringing scrutiny to bear on Wittgenstein's naturalistic concept of the "form of life," American sociologist Scott Lash takes into account the anthropological upheaval caused by the evolution of various mediums of thought on our technological forms of life (Lash 2002), a subject that has been abundantly discussed in the context of the Web (Halpin, Clark, and Wheeler 2014). Our main focus here, however, is less the future of humanity than that of philosophical research and philosophy itself. The architecture of the Web reveals a process of continuation and regrasping (which precisely needs to be properly assessed) of the most central of philosophical concepts: object, proper name, and ontology. On the Web, each concept of philosophy in its own way then gains a new existence as a technical artifact: objects turn into resources, proper names into URIs, ontology into Semantic Web ontologies.

Such a transition from philosophical concepts to technical objects isn't a one-way process and cannot remain without consequences for the original concepts that have been uprooted from their normal context, and accordingly this transition warrants careful examination. Do we philosophize today as we did in the past? With the same subject matter? Or in the same manner? Does it still make sense to locate oneself within established traditions, such as phenomenology and analytic philosophy, when their very own concepts freely cross these boundaries, and the real conversation is taken up elsewhere, using a language that only superficially seems identical to the one that preceded it? These kinds of questions have always been central to metaphilosophy, yet the advent of the Web—and so the philosophy of the Web—brings to these questions both a certain renewed importance and impetus. In the essays collected here, we bring together a number of authors who have offered some key contributions to this initial foray into the tentative realm of the philosophy of the Web. In order to guide philosophers through this nascent philosophical field, in the next section we delve deeper into the philosophical role of URIs and engineering as these two subjects serve as the twin foundations of the philosophy of the Web, and we then put each of the contributions in this collection within its philosophical context before reaching some tentative conclusions for next steps for the field.

1. URIs: "Artifactualization" of Proper Names

On the Web, the analogue of proper names is found in URIs, given by the standard IETF RFC 3986 to be "a simple and extensible means for identifying a resource," a definition in which resources are left crucially underdefined to be "whatever might be identified by a URI" (Berners-Lee,

Fielding, and Masinter 2005). URIs are everywhere: everything from `mailto:harry@w3.org` (for identifying an e-mail address of Harry Halpin) to `http://whitehouse.gov` (for identifying the page about the White House) qualifies as a URI. What quickly becomes apparent is that URIs are kinds of proper names for objects on the Web.

During the past fifteen years, philosophical discussions around the notion of a proper name have seamlessly followed in a business-as-usual manner, without any significant breakthrough. Yet during that same period, the architects of the Web have taken hold of the idea of proper names, and without purposefully altering its definition, have made naming the first supporting pillar of the Web, thus formulating an answer to the ages-old question of the relationship between words and things by combining in an original—and unintentional!—fashion the thoughts of Frege, Russell, Wittgenstein, and Kripke. For philosophy to take the URI, an engineered system of universal and accessible names, as a first-class philosophical citizen is then the first task of the philosophy of the Web.

While at first URIs may seem to be just a naming system for ordinary objects on the Web like e-mail addresses and web pages, the plan of Berners-Lee is to extend URIs as a naming scheme not just for the Web but for all reality—the Semantic Web will allow URIs to refer to literally anything, as “human beings, corporations, and bound books in a library can also be resources” (Berners-Lee, Fielding, and Masinter 2005). This totalizing vision of the Web is not without its own problems. In a striking debate between Berners-Lee and the well-known artificial intelligence researcher Patrick Hayes over URIs and their capacity to uniquely “identify” resources beyond web pages, Berners-Lee held that engineers decide how the protocol should work and that these decisions should determine the constraints of reference and identity, while Hayes replied that names have their possible referents determined only as traditionally understood by formal semantics, which he held engineers could not change but only had to obey (Halpin 2011). This duality can be interpreted as an opposition between a material and a formal a priori. Interestingly enough, recently, Hayes and other logicians such as Menzel have begun focusing on adopting principles from the Web into logical semantics itself, creating new kinds of logic for the Web (Menzel 2011). Unlike philosophical systems that reflect on the constraints of the world, the Web is a world-wide embodied technical artifact that therefore creates a whole new set of constraints. We suggest that they should be understood as a material a priori—in the Husserlian sense—grounded in history and technology.

Thus the Web, when it comes to its standards, breaks free from French philosopher Jules Vuillemin’s definition of a philosophical system as built on the logical contradictions between major philosophical schools of thought (Vuillemin 2009). Yet the Web doesn’t lead either to the collapse

of the transcendental and the ontological into the empirical, a new kind of “technological monism” as suggested by Lash (2002). *Logical contradiction* is overcome not by *factual opposition* (two words that Vuillemin highlighted) but through an *artifactual composition*, associating through the mediation of the artifact the virtues of competing philosophical positions. As the functions of concepts become functionalities, it is becoming increasingly easier to make them coexist for the sake of a *tertium datur*, without having to give up on consistency (Sloterdijk 2001).

The material a priori of technical systems such as the Web is brought about by what we call “artifactualization” (Monnin 2009), a process where concepts become “embodied” in materiality—with lasting consequences, as the result trumps every expectation, being more than a mere projection of preexisting concepts (which would simply negate the minute details of the object considered). While such a process clearly predates the Web, we can from our present moment see within a single human lifetime the increasing speed at which it is taking place, and through which technical categories (often rooted in philosophical ones) are becoming increasingly dominant over their previously unquestioned “natural” and “logical” counterparts. At the same time, the process of having philosophical ideas take a concrete form via technology lends to them often radically new characteristics, transforming these very concepts in the process. Heidegger posited a filiation between technology and metaphysics, with technology realizing the Western metaphysical project by virtue of technology inscribing its categories directly into concrete matter. Yet if technology is grounded in metaphysics, it is not the result of a metaphysical movement or “destiny” (*Schicksals*), but a more mundane contingent historical process, full of surprises and novelties. For all these reasons, it must be acknowledged that the genealogy of the Web, as a digital information system, differs from traditional computation with regard both to the concepts at stake and to our relation to them. The scientific *ethos* is indeed being replaced by an engineering one, something Berners-Lee dubbed “philosophical engineering” (Halpin 2008)—and this difference even holds true with regard to the (mainly logical, thanks to the Curry-Howard correspondence) *ethos* of computer science itself.

As already mentioned, URIs form the principal pillar of Web architecture, so it shouldn’t be surprising that they also constitute our gateway into the aforementioned problematic between engineering and philosophy. From its inception, the Web was conceived as a space of names, or “namespace,” even if the historical journey to URIs led through a veritable waltz of hesitations as the engineers who built the Web tried to pin down standardized definitions to various naming schemes. The numerous Web and Internet standards around various kinds of names bear witness to that ambivalence: URL (Uniform Resource Locator), URN (Uniform Resource Name), and even URC (Uniform Resource Characteristic or Citation). Each of these acronyms matches a different conception of the

Web and modifies the way it constitutes a system. Eventually, the acronyms have slowly evolved over time to return to Berners-Lee's original vision of a URI: a "Universal" Resource Identifier for everything, from which follows naturally the ability of links to allow everything to be interconnected on the Web.

The notion of proper name as it prevails today is directly inherited from analytic philosophy, and more precisely, from the work of Saul Kripke; although other definitions may exist in philosophy, Ruth Barcan Marcus (Humphreys and Fetzer 1998) is clearly the one who launched this Kripkean tradition, and this strand of work eventually meant that "proper name" would become the key operating term for questions on reference, identity, and modality. It holds such a weight that it explains how fields as diverse as epistemology and ontology can be considered part of a larger story, that of a *science of reference*. This space of convergence was historically opened by the different theories of intentionality and objects from Brentano to Twardowski and Meinong, but it was to split post-Frege philosophy into two rival traditions, the analytical and the phenomenological—the latter sometimes considered "continental" from the analytical perspective (Benoist 2001). Ruth Barcan Marcus and Saul Kripke's works on proper names provide the apex for the analytic tradition, but what we see now on the Web is the URI as a proper name and technical object that reopens a space for reunification of these two divergent philosophical traditions within the philosophy of the Web as the problems around reference and naming migrate from philosophical systems (Vuillemin 2009)—in particular, the philosophy of language—toward *technical* and *artifactual* systems, asking for a complete shift of analysis.

A clear example of how URIs are transforming the analytic tradition's bedrock of logic has recently been pioneered by Patrick Hayes, known for his original quest to formalize common-sense knowledge in terms of first-order logic of artificial intelligence but also more recently deeply involved for several years in the development of the Semantic Web, the extension of the Web beyond documents into a generic knowledge representation language (Hayes 1979). As a logical foundation for the Semantic Web, Hayes has suggested the creation of "Blogic" (a contraction of "Web logic," inspired by a similar contraction of "Web logs" to "blogs"), in which logical proper names, which possess no signification of their own outside their formally defined role in logic, would be replaced by dereferenceable URIs, which could in turn dereference logical sentences or even new interpretation functions not present in their original context.¹ Blogic would leverage the ability to use a name—in this case, a URI—to retrieve a "document," functionality that has played a critical part in the Web's

¹ See the talk "Blogic or Now What's in a Link?" by Patrick Hayes, online at http://videolectures.net/iswc09_hayes_blogic/.

success to this day, but outside hypertext and in the realm of semantics. With access mechanisms then possibly defining the semantics of proper names, the notion of reference on the Web cannot clearly choose between Wittgenstein (meaning determined by use), Russell (definite descriptions), and Kripke (rigid designators) for a theory of meaning. As a framework, the architecture of the Web composes with these conceptual positions: a user is free to give any kind of meaning to a URI, someone publishing a new URI may refer to it rigidly or with the help of a description for the Semantic Web, and what we access via that URI can also play a role in defining its meaning.

This mixture of the technical and the philosophical is found not only in the semantics of URIs but also in their governance (the latter having an impact on the former). URIs are not just free-floating names but assigned virtual territory controlled by bodies such as the Internet Assigned Numbers Authority (IANA) via domain registrars. While being proper names, URIs also have a particular legal and commercial status that does not clearly compare to proper names in philosophy, with perhaps only a vague analogue to the ability of organizations to copyright names. For websites, controlling a name, and thus which web pages can be accessed from it, is a source of immense power. According to Tim Berners-Lee, the ability to mint new URIs and link them with any other URI constitutes not only an essential linguistic function but also a fundamental freedom. Yet as URIs leave the field of semiotics, they undergo a change in nature as regards both the possibilities offered by a technology of naming and the limitations imposed by the legislation that governs bodies such as IANA. Once again, objects such as URIs or disciplines such as philosophy that seemed purely formal are gaining a newly found materiality, full of historical and even political contingencies around extremely concrete juridical and economic issues.

2. Denaturalizing Ontology: Philosophical Activity Redux

Another field in which the Web is rapidly causing massive conceptual tremors is the once forgotten philosophical realm of ontology. Given the long-lasting gap between a name and its object, the study of URIs on the Web naturally causes an intrepid philosopher of the Web to lean on work on names in philosophy of language, while with the study of ontology on the Web we return to the preponderance of the object. Generally considered a branch of metaphysics, ontology traditionally has generally been the study of the (often possible) existence of objects and their fundamental categorization and distinctions. Interestingly enough, engineering-inclined artificial intelligence researchers (the late John McCarthy being first among them) have also seized upon the word “ontology” over the past fifty years, making “ontology” their own term for purposes of creating

knowledge representation languages, as exemplified by Gruber's famous engineering definition that "an *ontology* is a specification of a conceptualization" (Gruber 1993). While this definition may at first glance seem so vague to be totally useless, one should remember that Tim Berners-Lee also had notorious trouble defining precisely what a URI is, and this did not seem to prevent URIs from becoming central to the entire edifice of the Web. In fact, one would almost suspect that the utility of a term may somehow be related to the fact that it is underdefined—or perhaps more precisely, defined just enough to allow concrete engineering to reveal the inherent productivity of the term's concept. While the use of ontology by knowledge representation has become sidelined in philosophical circles by more clearly philosophical debates in artificial intelligence around embodiment, the move of Berners-Lee to create a Semantic Web that transforms the Web from a space of URIs for hypertext documents to a giant global knowledge representation language built on URIs has led to a renewal of interest in the engineering of ontologies as well. We suspect this computational turn in ontologies on the Web will in turn lead to a revival of the philosophical field of ontology.

As the shift from philosophical ontology to ontological engineering progresses, philosophers are gradually losing control over their own tools, even if they are not necessarily aware of it. What ensues is a real "proletarianization," as Bernard Stiegler (1998) puts it, and this process is smooth and passive, since philosophical activity goes on uninterrupted, as if nothing were amiss. Nonetheless, there are a number of unmistakable signs. Following the example of Barry Smith, some philosophers have already made their move explicit by rebranding themselves ontologists, as they are now working exclusively in the field of knowledge engineering. Might the conundrum of this technological life form be all about employing the concepts of philosophy in a new light while at the same time making the previous blissfully technologically unaware philosophical discipline obsolete? Although there are possibly some examples to illustrate such a strong point, a more reasonable response would be to answer this question with a little more subtlety by taking into account the precise nature of how the Web transforms ontology before tackling the wider question of how the Web transforms philosophy.

The ontological implications of the Web are deeply related to the concept of a resource harbored at the heart of Web architecture; for the philosophy of the Web this particular concept constitutes an opportunity to renew the question of ontology itself. As designated by a URI, a resource can be "anything at all," exactly as was the case for the hoary philosophical concept of the "object," which was the actual focus of the ontological tradition (as long as you trace the word back to its origins in the seventeenth century, more than twenty centuries after Aristotle's definition of the science of being). Consequently, it is not the sole business of philosophers, hidden far from the world at the back of some unidentified

abode where they can hone their weapons alone to issue judgments and decide what “anything at all” really means. Everywhere, the gap is narrowing and the old privileges are in crisis, as Scott Lash (2002) has discussed. The new indexing tools and the contributive nature of the Web make it possible for anyone to tackle this issue—not only philosophers and engineers. Thus, the way in which the question will be asked relies on the Web itself. Nevertheless, in opposition to Lash’s thesis, ontology is not suffering from being conditioned by the patterns of our technological way of living. Technology is the condition of the liberation of ontology, the “ontogonic” dimension of technology discussed by Bruno Bachimont (2010) and more recently with regard to the notion of philosophical engineering in Monnin (2012b) that draws on Pierre Livet’s work presented in this collection (Livet 2014).

According to Livet’s recent work on the ontology of the Web, we can chart the operations that allowed the emergence of the objects that in turn proved essential to conceptualize the Web’s architecture and, subsequently, to clear the ontological horizon (Livet 2014). Far from consisting merely of epistemic processes, this work opens the door for an ontology of operations constitutive of an ontology of entities, which bit by bit refines itself as time goes by. The possibility to move back and forth, as the whole process unfurls, is not to be excluded, leading to new beginnings and thus leaving entirely open the question of the nature of the ultimate constituents of our technological cosmos.

Given the conceptual purity required by formal ontology, built as it is upon relationships of dependency and the application of mereology, and the materiality of devices as the place of a new technicized a priori, the time has come for a re-evaluation of the very notions of form and matter, through the filter of digital technologies in general and the Web in particular. In this regard, the research initiated several years ago on *Ontology Design Patterns* (Gangemi and Presutti 2009), which may appear to be limited only to the field of knowledge engineering, in fact has implications far beyond the boundaries of its original field. Incidentally, nothing prevents philosophers from trying to conceive their practice in a more collaborative fashion using a similar pattern analysis of their own activity, in order to gain a better view on the collective fine-grained ontological invariants that groups of philosophers share beyond the explicit debates through which philosophers normally distinguish themselves. Beyond this, the background against which these ontological patterns appear is formed by practices that, though they may produce certain apparently transhistorical regularities, are rooted in an historical context and therefore should not be “naturalized” prematurely.

In order to identify and qualify these invariants by taking into account that which supports and maintains them, one has to, so to speak, “denaturalize ontology”—and this slogan could serve as a synthesis of the entire philosophical research program we are suggesting here. To pretend, as is

often the case in analytic philosophy, that certain ontological constructions are simply pre-given would be a serious error, for everything has a cost—one need only consider the works of Bruno Latour and Pierre Livet to realize that (Latour 2001; Livet 2014). The key is determining how technology opens an avenue into the historicization of ontology.

There are clear predecessors in either explicitly or implicitly building ontologies into technology as well as having technology influence our everyday ontology. In the fields of cognitive science and artificial intelligence, the question of the representation, formalization, and computation of knowledge—as well as the more philosophically neglected approaches centered on collective intelligence and human computation that partially go beyond traditional philosophy of the mind due to their focus on human intelligence’s complementarity with the machine—have already produced interesting leads in the wake of the work of Andy Clark and David Chalmers on what they call “the Extended Mind Hypothesis” (Clark and Chalmers 1998). We would like to extend the extended mind by renewing metaphysics through a focus on the positive aspect of French linguist and semiotician François Rastier’s critique of cognitive science, which he accused of “naturalizing metaphysics” (Rastier 2001).² Given the lessons learned from the implicit metaphysics of cognitive science and artificial intelligence, we cannot simply criticize or reject the ongoing exploitation of the Web on an unprecedented scale that harnesses vast resources of centuries of philosophical debates on language and knowledge. In order to describe the paramount importance of the technical production of media ranging from television to computers in the twentieth century, Bernard Stiegler (1998) coined the expression “the machinic turning-point of sensibility.” With a slight shift of focus, we may talk of a *machinic* (or perhaps better, *artifactual*) *turning-point of metaphysics* itself, an ongoing deep modification of the meaning of metaphysics in philosophy.

At the present moment, the perspective given by the original architecture of the Web needs to be broadened as two problematics currently intersect: (a) the *artificialization* of a growing number of particular domains (“natural” but also “formal” ones,³ each of these two notions being traditionally contrasted with technics but now becoming technical) and (b) the *artifactualization* of philosophical concepts in general as they are imported into the realm of the digital—in particular under the guise of

² Of course, one of the necessary conditions to be spared the risk of unduly naturalizing metaphysics is to make sure that no kind of prophetic and biologically inclined conception of technology is allowed to thrive in parallel.

³ The origin of this paradox could be found in Husserl’s work, in which formal and material a priori and ontologies are articulated. As always, this is only a starting point, and this distinction is likely to be taken up and revised by other currents of thought. We think about material cultures or authors whose works, although they differ in many ways, all put the stress on the materiality of the mediums and operations of knowledge.

the dominant sociotechnical form of the twenty-first century, which is none other than the Web. It may even be argued that semiotic objects, such as philosophical concepts, already have features similar to those of technical objects (Halpin 2008),⁴ in which case this latest round of digitization on the Web may rather appear tantamount to a *re-artifactualization*, provided that we do not ignore the (often overlooked) original ties of philosophy to technology.

Behind the distinction between the two aforementioned problematics stands an important issue: if these two dimensions are not clearly acknowledged, there is a risk that we will “naturalize” (a) without any real philosophical scrutiny the re-introduction (b) of some body of philosophical knowledge (or some unconscious philosophical legacy) while designing technical systems.⁵ For all these reasons, the very practice of philosophy is transformed by having to take the material a priori and its technical categories as seriously as “natural” (synthetic) or “analytic” categories from biology or natural language. Philosophers then have to deal with engineered categories that may have a lasting effect in domains like the Web, not just as variants of categories that can be analytically understood but rather as concrete artifacts that can even transform analytic categories previously taken for granted. Ironically, the main challenge to analytic judgment is no longer what Quine called naturalization, but rather the ongoing artifactualization of which the Web is the historical exemplar par excellence (Livet 2014; Monnin 2013).

3. Open Problems of the Philosophy of the Web

Now that we have surveyed some of the core foundations of the nascent philosophy of the Web, we should be clear that we have only started to embark on this particular route, and the road ahead lies littered with open philosophical problems worthy of a tassel of theses. The Web was the brainchild not of Tim Berners-Lee as a lone individual but of a large and heterogeneous group of Web architects, ranging from Berners-Lee’s compatriots involved in standards, such as Larry Masinter and Roy Fielding (Berners-Lee, Fielding, and Masinter 2005), to users who contribute

⁴ On the articulation of semiotic and technical aspects, see also the work of Bruno Bachimont (2010).

⁵ In this view, the warning is also useful against all ways of thinking that proceed step by step (two steps, to be precise), “computationalizing” or “informationalizing” the world first, in order to be subsequently entirely free to naturalize ontology, which is afterward considered a natural science of the universe (a tendency that historically dates back to the Neoplatonic philosophers). Opposing this way of thinking, Jean-Gabriel Ganascia advocates a computerized epistemology that is directly related to the cultural sciences, an epistemology that is well aware of the status of computers in the production of contemporary knowledge. In this respect, our point of view is quite close to his (Ganascia 2008).

content daily to the Web; the philosophy of the Web must likewise be a collective affair. We will provide brief summaries of some of the most pressing questions that face the philosophy of the Web, along with an overview of those contributors in this collection who have addressed them.

3.1. What Is the Relationship of the Philosophy of the Web to a More General Philosophy?

It goes almost without saying that every new discipline must build its foundations on earlier philosophical studies and situate itself consciously within a wider historical context, yet to do this correctly is one of the most difficult tasks for a philosopher. Case in point: one of the factors distinguishing the Web from earlier visions of an interconnected global information space, such as Vannevar Bush's Memex (Bush 1945), is that the Web is implemented on digital computers. This simple decision to stick to a digital medium had a deep impact on engineering matters, allowing as it does functionality from the copying and caching of pages used by search engines to the transformation of music and video by streaming media. Yet this almost obvious engineering decision in turn leads to a decidedly deep impact on metaphysical notions such as that of a resource (Berners-Lee, Fielding, and Masinter 2005). Yet in expanding the Web into the Semantic Web, where seemingly non-digital things are construed as resources, there still seems something self-evidently different between web pages and things themselves, as one can easily copy pages about the Eiffel Tower, due to their being digital, but not the Eiffel Tower itself. However, precisely explicating the philosophical difference between a page about the Eiffel Tower and the Eiffel Tower itself seems to always involve punting the question to an under-theorized notion of digitality. While there has been considerable philosophical debate over the nature of logic and computation, there has been little work on the wider notion of digitality. One of the prime tasks of the philosophy of the Web is to determine how the Web is engineered on top of robustly digital objects. This is not to say that the Web must remain digital forever—as the Web transforms into an Internet of Things and increasingly interacts with the analogue world, understanding digitality becomes more—not less!—important, as does situating this digital turn within wider currents in the analytic and phenomenological traditions. Yuk Hui's contribution "What Is a Digital Object?" provides a synopsis of his much larger foray into this field, and artfully combines Husserl's phenomenology with Simondon's understanding of technics to provide just such a theoretical foundation for the philosophy of the Web (Hui 2014).

3.2. Does the Web Radically Impact Metaphysics, Ontology, and Epistemology?

As we indicated earlier, it appears that the Web is having perhaps its most crucial role in philosophical realms that seem at first glance rather

distant from high-speed technology. Taking on board the efforts of the Semantic Web to rework the logical foundations of ontology, Monnin (2014) argues that the main innovation of the architecture of the Web is definitely more ontological than technical. Raising the question “What do Web identifiers refer to and how?” he examines the answer provided by Web architects themselves. The investigation, in a way, is reminiscent of Quine’s landmark paper “On What There Is,” once it has been brought to the Web. Monnin’s conclusion is that the Web articulates a deeply subtle view of objects on a global scale, theoretically as well as technically. In other words: the Web is an operative “ontology,” as suggested in the title of his piece. Pierre Livet demonstrates in his contribution “Web Ontologies as Renewal of Classical Philosophical Ontology” how far the Web has taken us from classical ontological questions based on “natural” kinds to new dynamic and open-ended ontologies (Livet 2014). On a similar note, without a doubt the Web seems to be impacting the phenomenology of such fundamental metaphysical categories as space and time for ordinary users. Michalis Vafopoulos’s contribution “Being, Space, and Time on the Web” is precisely such a retheorization of metaphysics on the Web, taking such fundamental aspects of the Web as the number of links in a page and reconceiving of this as a concept of space, with the time spent by users visiting a given resource as a concept of time, and then drawing a number of social and economic conclusions (Vafopoulos 2014). Moving from the world itself to our knowledge of the world, we find that one of most interesting phenomena brought about by the Web is the tendency of people to increasingly rely on search engines to answer their everyday questions. Precisely how ubiquitously search engine usage impacts classical conceptions of epistemological questions of knowledge and belief are tackled in the piece by Thomas Simpson, “Evaluating Google as an Epistemic Tool” (Simpson 2014). These epistemological questions are also not purely theoretical, as obviously the use of search engines in everyday environments, political debate in wikis and (micro)blogs, and the possible transformation of the university system itself by massive Web-mediated online courses all merit serious attention, and how we understand these issues is directly influenced by our position on whether or not access to knowledge on the Web counts as belief or even knowledge.

3.3. Can Human Cognition and Intelligence Genuinely Be Extended by the Web?

Questions about epistemology naturally lead to questions about whether or not the Web is changing our conception of humanity, and to questions about how the philosophy of the Web interacts with other empirically informed philosophical questions around neuroscience and cognitive science. In particular, the question of whether or not human cognition is genuinely extended by the Web appears rather naturally: In

a not-so-distant common scenario when humans are wearing Google-enabled goggles that allow them to almost instantly and seamlessly access the Web, would one give the Web some of the cognitive credit for problem solving? This very example, in terms of a special pair of glasses rotating blocks in the classic game of Tetris, is brought up as one of the motivating examples of the classic Extended Mind Hypothesis of Chalmers and Clark (Clark and Chalmers 1998), and now this example is coming very close to reality on the Web. However, the precise conditions of what constitutes the mark of the cognitive, and under what conditions the Web actually counts as part of an extended mind, are still very much a topic of debate. The first essay in this collection, “Philosophy of the Web: Representation, Enaction, Collective Intelligence,” outlines connections between the philosophy of the Web and what has been termed “‘4E’ (embodied, embedded, enactive, extended) cognition (Halpin, Clark, and Wheeler 2014). In this overview, Halpin, Clark, and Wheeler note how concepts from cognitive science such as representation and enaction must be re-thought through in the light of the Web’s status as a readily accessible externalized public cognitive resource. They end with broaching the topic of how the Web brings forth the possibility of not just an extended mind (Clark and Chalmers 1998) but also a new type of massively distributed collective intelligence that has yet to be properly studied via the lens of philosophy. In his contribution, “The Web-Extended Mind,” Paul Smart mounts an argument for why the Web should be considered part of the extended mind, and then proceeds to show how engineering developments such as Berners-Lee’s Semantic Web may have certain design characteristics that could more tightly bind the mind to the Web in the future (Smart 2014). This naturally implies revisiting the question of intelligence. Interestingly enough, in the early days of the Internet, the motivating vision was one of collective intelligence, in which machines augmented rather than replaced human intelligence. This vision was a direct rival of artificial intelligence, which hoped to implement human-level intelligence in machines. Surprisingly, a thorough investigation of the philosophical assumptions and differences of both collective and artificial intelligence has yet to be written, an investigation that has become increasingly urgent, as the Semantic Web itself is often criticized as a mere repeat of classical artificial intelligence. A first important step has been taken by Selmer Bringsjord and Naveen Sundar Govindarajulu in their contribution “Given the Web, What Is Intelligence, Really?” in which they convincingly argue that even if the Semantic Web did become a reality, it would lack the reasoning capacity of humans that intelligence requires (Bringsjord and Govindarajulu 2014). How this particular debate over the transformation of intelligence on the Web plays out in the future will no doubt have yet unforeseen ramifications, just as the original quest for artificial intelligence radically revised the traditional pre-computational philosophy of the mind.

3.4. Does the Web Alter Our Domain-Specific Practices in a Manner That Demands a New Qualitative Analysis?

Given our order of presentation, one of the central questions of the philosophy of the Web would seem to be how the new engineering-inspired revolution in the rarefied air of philosophy will have its tremors felt in various specialized domains. Quite the reverse seems true: the impact of the Web has most powerfully been noticed in its empirical effects on almost impossibly heterogeneous domains, ranging from online recommendation systems in e-commerce to the near instantaneous spread of news globally via microblogging. Indeed, the primary difficulty of the philosophy of the Web lies precisely in the difficulty inherent in tracing how such a diverse range of complex induced effects could form a coherent philosophical system, one that may have explanatory and even predictive power. Thus, detailed domain-specific studies of how the Web impacts particular domains of practice are critically part of the philosophy of the Web. Few areas can be considered seemingly more remote from mundane engineering considerations than proof-proving in mathematics, yet in their remarkable contribution, “The Web as a Tool for Proving,” Petros Stefaneas and Ioannis Vandoulakis demonstrate the nature of the radical impact of the Web on this most formal and theoretical of domains (Stefaneas and Vandoulakis 2014). To move in the reverse direction, the communication and ubiquitous accessibility of the Web may alter our notion of embodiment. Nowhere is this more powerfully demonstrated than in the multiplayer Web-mediated virtual worlds, whose numbers seem to be growing everyday. While currently only a small part of the Web, it is very possible that such deeply immersive and even “three-dimensional” environments may come of age soon and become an important part of the future of the Web. Johnny Hartz Soraker engages with these environments in “Virtual Worlds and Their Challenge to Philosophy: Understanding the ‘Intravirtual’ and the ‘Extravirtual,’” where he carefully compares and contrasts the kinds of actions possible in these worlds (Soraker 2014). This leads one to think that fundamentally the Cartesian distinction between the “real world” and the “virtual world” may indeed be far more complex than initially conceived, positing a problematic that may end up being just as important for the philosophy of the Web as the mind–body problem is for the philosophy of mind. Far from being purely academic, these debates over how the Web interacts with our daily life are already stirring upheaval in how we understand our own notion of privacy and identity, and will soon perhaps even take political center stage.

3.5. The Future of the Philosophy of the Web

In this collection, we have endeavoured to take into full account both the engineering aspect and the wider philosophical ramifications of the Web.

So we are pleased to feature new pieces from both Tim Berners-Lee and Bernard Stiegler on philosophical engineering and the philosophy of the Web. We are privileged to include with his permission the text of an interview we did with Berners-Lee, widely acclaimed as the inventor of the Web (Halpin and Monnin 2014). In this wide-ranging interview, Berners-Lee reveals why he coined the term “philosophical engineering” and offers his thoughts on the future of the philosophy of the Web. We end the collection with an afterword entitled “Web Philosophy” by Bernard Stiegler, one of France’s preeminent philosophers of technology, who uncovers the potential and the “shadows” at the heart of the digital enlightenment’s new political philosophy (Stiegler 2014). Stiegler attempts to place “philosophical engineering” in proper context by uncovering its lineage from Archimedes to the Web by way of philosophers such as Plato, Husserl, and Derrida.

4. Conclusion

Ultimately, the philosophy of the Web has just begun, and its future is far from certain: the impact of the Web may ultimately be as transformative as that of natural language, or perhaps it will be superseded within a short time by some truly distinct technological development. Regardless, just as the study of artificial intelligence provoked genuine philosophical inquiry into the nature of mind and intelligence, the philosophy of the Web will at least—we hope!—provoke the taking of engineered artifacts such as the Web that impinge on areas traditionally the province of philosophy as first-class subjects of concerted inquiry by philosophers. The Web has obviously benefitted from previous encounters with philosophically informed engineers, although this point alone would deserve a more ample treatment (Shadbolt 2007).

All these engineering-related activities on the Web can easily be described, if one is willing to alter Clausewitz’s adage, as “philosophy continued through other means.” Most of these means, though, are far from being completely exogenic. It is only through a redoubling of awareness over these new technological mediations, studied and developed by “philosophical engineering,” that philosophers may have the opportunity to extend their categories by submitting them to the evaluation of the non-human technical artifacts, now regular constituents of our world. The stake here is a change of course that has nothing to do with just switching one given a priori (or *epistēmē*) for another. Ian Hacking (2002) sharply underlined that such notions are far too massive. We believe that those interested in the philosophy of the Web are not simply wearing a new pair of conceptual glasses; the world itself has changed, for it is composed no longer of canonical entities but of utterly new ones that differ fundamentally from their predecessors. Enriched with the new details bequeathed to

it by technology, the world is asking for a redefinition of its ontological cartography (perhaps an “ontography”), even if it implies that we should broaden our philosophical focus to encompass all the agents responsible for the shift to the Web (engineers, languages, standardization committees, documents, search engines, policymakers, and the like). No empirical metaphysics is entitled to define the “nature,” patterns, or limits of the agents that can impact the philosophy of the Web without argumentation and clarity. What is at stake, the determination of the “collective,” the cosmos we live in, requires that we stay on the verge of philosophy itself, on the very technical spot where new objects⁶ are spreading, and once these objects are brought back to their philosophical womb, they will certainly foster mutation within philosophy. We consider the Web to be endowed with this capacity more than any other technical apparatus, and that the duty of the philosophy of the Web is to set ourselves upon the task.

As we have seen, this task leads to a number of open questions, and we hope that the contributions to this collection have put forward some of these questions as clearly as possible, in order to highlight theoretical as well as technological issues—and even social and political matters—that will set the philosophy of the Web in motion toward the interdisciplinary point of view necessary to adequately address these problems. In the context of the Web more than ever, neither philosophy nor engineering can escape its practical consequences by dodging issues that are relevant, and even vital, to how the architects of the Web have in the past engineered and will in the future engineer the Web itself, as it increasingly becomes the primary medium of knowledge and communication. Again, the traffic is not only one-way between philosophy and the Web, it is a dynamic feedback cycle: as the Web itself is mutating as a medium, we can consider humans and objects of knowledge to be condemned to mutate in turn. Historically, philosophy is a discipline descended from the alphabet and the book. Under the influence of what has been called by French historian of language Sylvain Auroux a “third revolution of grammatization” (Auroux 1994) catalyzed by digital technologies and the Web, what turn philosophy takes remains to be seen.

Richard Sennett’s motto in a recent book appears to be “doing is thinking” (Sennett 2008); once concepts have been artifactualized (and, as a consequence, externalized), thinking is also doing, and so in the end, a matter of design. In this regard, we need to reject Marx’s Eleventh Thesis on Feuerbach: in the era of the Web, interpreting the world is already changing it. This holds true especially when the art of interpretation and theorizing is serving the purpose of the creation of new forms of technology that harness the power of the Web via building on and creating

⁶ Several of these new objects have been minutely examined by other disciplines, and we may want to borrow from them. One meaningful example is the analysis of standards as an essential tool to understand the Web’s architecture.

Web standards. These innovations target a reality yet to come, but a reality that we can already conceive of, as opposed to a pre-existing reality. The matter of engineering the Web is rapidly transforming into the matter of engineering the world. To paraphrase Saussure, the philosopher and the engineer are both challenged to acquire a clear view of what it is they are doing. For the philosophy of the Web and philosophy itself, these are the stakes.

Afterword

A few words are necessary in order to provide some context for the recent development of the philosophy of the Web that led to this collection, and to provide a call for those interested to join in shaping the future of the philosophy of the Web. The term “philosophy of the Web” was first coined by Halpin (2008), while the central importance of artifactualization in the philosophy of the Web was first explicated by Monnin (2009). Monnin organized the first Web and Philosophy symposium in 2010 at La Sorbonne. Subsequent editions of this symposium were organized jointly by Halpin and Monnin in 2011 at Thessaloniki to coincide with the conference on the Theory and Philosophy of Artificial Intelligence and at Lyon in 2012 as a workshop at the International World Wide Web Conference. Several of the contributions to this collection are extended versions of presentations given at these events, and we would like to thank all of those who have participated in the discussion so far. Again, the philosophy of the Web is not the static product of a single individual (or even two!), but a collective endeavour like the Web itself, whose scope and power widens the more that people are involved. To engage in future events and public debate, join the W3C Philosophy of the Web Community Group: <http://www.w3.org/community/philoweb/>.

References

- Auroux, Sylvain. 1994. *La révolution technologique de la grammatisation*. Liège: Mardaga.
- Bachimont, Bruno. 2010. *Le sens de la technique: Le numérique et le calcul*. Paris: Encre Marine.
- Benoist, Jocelyn. 2001. *Représentations sans objet: Aux origines de la phénoménologie et de la philosophie analytique*. Paris: P.U.F.
- Berners-Lee, Tim, Roy Fielding, and Larry Masinter. 2005. “IETF RFC 3986—Uniform Resource Identifier (URI): Generic Syntax” Available at <https://tools.ietf.org/html/rfc3986>.
- Bringsjord, Selmer, and Naveen Sundar Govindarajulu. 2014. “Given the Web, What Is Intelligence, Really?” Included in this collection.

- Bush, Vannevar. 1945. "As We May Think." *Atlantic Magazine*, vol. 7. Available at <http://www.theatlantic.com/magazine/toc/1945/07>.
- Clark, Andy, and David Chalmers. 1998. "The Extended Mind." *Analysis* 58:10–23.
- Ganascia, Jean. 2008. "'In Silico' Experiments: Towards a Computerized Epistemology." *APA Newsletter on Computers and Philosophy* 7, no. 2:11–15.
- Gangemi, Aldo, and Valentina Presutti. 2009. "Ontology Design Patterns." In *Handbook on Ontologies*, second edition, edited by S. Staab and R. Studer, 221–43. Berlin: Springer.
- Gruber, Thomas. 1993. "A Translation Approach to Portable Ontology Specifications." Available at <http://tomgruber.org/writing/ontolingua-kaj-1993.htm>.
- Hacking, Ian. 2002. *Historical Ontology*. Cambridge, Mass.: Harvard University Press.
- Halpin, Harry. 2008. "Philosophical Engineering: Towards a Philosophy of the Web." *APA Newsletter on Philosophy and Computers* 7, no. 2:5–11.
- . 2011. "Sense and Reference on the Web." *Minds and Machines* 21, no. 2:153–78.
- Halpin, Harry, Andy Clark, and Michael Wheeler. 2014. "Philosophy of the Web: Representation, Enaction, Collective Intelligence." Included in this collection.
- Halpin, Harry, and Alexandre Monnin. 2014. "Interview with Tim Berners-Lee." Included in this collection.
- Hayes, Patrick. 1979. "The Naive Physics Manifesto." In *Expert Systems in the Micro-Electronic Age*, edited by D. Michie, 242–70. Edinburgh: Edinburgh University Press.
- Hui, Yuk. 2014. "What Is a Digital Object?" Included in this collection.
- Humphreys, Paul, and James Fetzer. 1998. *The New Theory of Reference: Kripke, Marcus, and Its Origins*. Berlin: Springer.
- Lash, Scott. 2002. *Critique of Information*. London: Sage.
- Latour, Bruno. 2001. *Pasteur: Guerre et paix des microbes, suivi de Irréductions*. Paris: La Découverte.
- Livet, Pierre. 2014. "Web Ontologies as Renewal of Classical Ontology." Included in this collection.
- Menzel, Christopher. 2011. "Knowledge Representation, the World Wide Web, and the Evolution of Logic." *Synthese* 182, no. 2:269–95.
- Monnin, Alexandre. 2009. "Artifactualization: Introducing a New Concept." In *Proceedings of Interface09, 1st National Symposium for Humanities and Technology*. Southampton: University of Southampton.
- . 2012a. "L'ingénierie philosophique comme design ontologique: Retour sur l'émergence de la 'ressource.'" *Réel-Virtuel* 3, Archéologie

- des nouvelles technologies. Available at <http://reelvirtuel.univ-paris1.fr/index.php?/revue-en-ligne/3-monnin/2/>.
- . 2012b. “The Artifacts of Reference and ‘Substances’ on the Web: Why (HTTP) URIs Do Not (Always) Refer nor Resources Hold by Themselves.” *American Philosophical Association Newsletter on Philosophy and Computers* 11, no. 2:11–9.
- . 2014. “The Web as Ontology.” Included in this collection.
- Rastier, François. 2001. *Sémantique et recherches cognitives*. Paris: P.U.F.
- Sennett, Richard. 2008. *The Craftsman*. New Haven: Yale University Press.
- Shadbolt, Nigel. 2007. “Philosophical Engineering.” *Words and Intelligence II: Essays in Honor of Yorik Wilks*, edited by Khurshid Ahmad, Christopher Brewster, and Mark Stevenson, 195–207, Text, Speech, and Language Technology series 36. Dordrecht: Springer.
- Simpson, Thomas W. 2014. “Evaluating Google as an Epistemic Tool.” Included in this collection.
- Sloterdijk, Peter. 2001. “Die Domestikation des Seins: Für eine Verdeutlichung der Lichtung.” In *Nicht gerettet: Versuche nach Heidegger*, edited by P. Sloterdijk, 142–234. Berlin: Suhrkamp.
- Smart, Paul R. 2014. “The Web-Extended Mind.” Included in this collection.
- Søraker, Johnny Hartz. 2014. “Virtual Worlds and Their Challenge to Philosophy: Understanding the ‘Intravirtual’ and the ‘Extravirtual.’” Included in this collection.
- Stefaneas, Petros, and Ioannis M. Vandoulakis. 2014. “The Web as a Tool for Proving.” Included in this collection.
- Stiegler, Bernard. 1998. *Time and Technics: The Fault of Epimetheus*, Volume 1. Palo Alto, Calif.: Stanford University Press.
- . 2014. “Afterword: Web Philosophy.” Included in this collection.
- Vafopoulos, Michalis. 2014. “Being, Space, and Time on the Web.” Included in this collection.
- Vuillemin, Jules. 2009. *What Are Philosophical Systems?* Cambridge: Cambridge University Press.