The decentralization of knowledge: How Carnap and Heidegger influenced the Web by Harry Halpin and Alexandre Monnin

Abstract

Does the centralization of the Web change both the diffusion of knowledge and the philosophical definition of knowledge itself? By exploring the origins of the Semantic Web in the philosophy of Carnap and of Google's machine learning approach in Heidegger, we demonstrate that competing philosophical schools are deeply embedded in artificial intelligence and its evolution in the Web. Finally, we conclude that a decentralized approach to knowledge is necessary in order to bring the Web to its full potential as a project for the spread of human autonomy.

Contents

[Introduction](#)

[The original \(decentralized\) vision of the Web](#)

[Carnap and the philosophical roots of knowledge representation and the Semantic Web](#)

[From Heidegger to knowledge graphs and machine learning](#)

[Conclusion](#)

Introduction

Can we decentralize the Web given its current advanced state of centralization into “walled gardens” controlled by a few large platforms? The premise inherent in this question is that at some primordial point the Web was decentralized and, in proper technological millenarian form, the Web can be decentralized yet again. Yet, often left unstated in this proposition is that at its heart, the fate of decentralization goes far beyond the Web: The future of planetary knowledge is at stake. We argue that this is not just a political debate over the control of knowledge, but a philosophical debate over the nature of knowledge itself.

In other words, the danger of centralization is not only an existential threat to the open Web, but also to the larger philosophical project that underwrites the existence of the Web as a knowledge infrastructure. The increased centralization of the Web is inarguable, as only two companies — Facebook and Google — control more than half the flow of traffic throughout the Web in 2016. However, such centralization is not predestined nor the result of a conspiracy; a more sound argument for the centralization of the Web is that centralization is fundamentally structural to any maturing industry. All maturing capitalist industries eventually become oligopolies, so that the centralization of the Internet into a few increasingly feudal fiefdoms simply shows the Web cannot escape the same pattern of classical pre-Internet telecommunications and automobile industries (Wu, 2011). However, we will argue that the Web is not just another industry, but possesses a special epistemic import as the latest incarnation of a larger progressive philosophical project of the decentralization of knowledge, a philosophical project to ultimately advance human autonomy.

The Web is not merely technology, but philosophical ideas given technological flesh: The Web is the latest incarnation of the Enlightenment project to renew the promise of philosophy for self-knowledge, but this time as a “digital enlightenment” (Bus and Crompton, 2012). Technology may be seen here as the continuation of philosophy by other means (and for different aims), and so the Web can be thought of as a battleground between different conceptions of knowledge, ranging from the logical empiricism of Carnap to the focus on embodiment of Heidegger, all of which are reflected in the epistemological and technical underpinnings of software as diverse as Wikipedia and deep learning. We must be wary as well, as the philosophical hypothesis that the Web is part of the Enlightenment project is itself far from unproblematic: There is a distinct possibility that the Web itself has also led to new and more potent forms of domination, exploitation, and inequality — as with Enlightenment, there is a dark side to the Web. Our goal is to retrieve the promise of the Web as a platform for knowledge, as latent in the

promise of the re-decentralization of the Web is that the Web itself is not fated to merely repeat the mistakes of the Enlightenment but that some form of genuine autonomy is globally possible due to the spread of knowledge, no longer confined to an elite minority.

This philosophical project of epistemological autonomy is profoundly political to its core, far beyond both the typical “hype” from Silicon Valley about changing the world via some short-lived gadget. Indeed, what we may be seeing with the centralization of the Web is not just a maturing of yet another industry, but also the fundamental closing off of possibilities — technical, political and organizational — that are necessary for confronting the myriad social, economic, and ecological crises that lie at the heart of the Anthropocene — now and for decades to come (Latour, 2014). Thus, we will not dwell on the precise economic arguments for either the original success of the Web or its eventual centralization, nor on the economic debates around its possible re-decentralization. We will likewise leave aside debates over the best technical or political means to re-decentralize the Web. Instead, we will address the philosophical history of the Web, and how this philosophy is in need of renewal in order to help answer the above pressing matters.

First, we will argue that the decentralization of knowledge was built into the architectural design of the original Internet and Web, and that this intention led to its world-historical success: The wiring of over three billion people into a single global epistemological environment. The Web is for more than just cat memes; The spread of knowledge can be evidenced by the well-trodden example of Wikipedia, where Diderot’s original vision of a universal encyclopedia of all human knowledge has been resurrected for a digital age and is now increasingly globally accessible. However, the next step in the evolution of knowledge as foreseen by the inventor of the Web, Tim Berners-Lee — the Semantic Web as a universal space of data (Berners-Lee, *et al.*, 2001) — never truly materialized, despite being rooted in the philosophy of Carnap as explored in traditional artificial intelligence and knowledge engineering (Monnin, 2015). What did emerge was a number of proprietary “knowledge graphs” that harvested the production of knowledge inherent to Wikipedia.

The alternative to the Semantic Web that has emerged is massive machine learning, which now, in the form of “deep learning”, is increasingly tackling the semantics once thought to be the exclusive province of human-produced knowledge representations. In combination with the open data produced by the Web, knowledge graphs and deep learning algorithms serve as the motor behind Google, Facebook and the current centralization of the Web. Similar to artificial intelligence, machine learning is also a philosophical project, with its own non-conceptual theory of knowledge that can be traced to the pragmatic reading of Heidegger given by Winograd, thesis advisor of Larry Page, CEO of Google (Winograd and Flores, 1986).

In this regard, the struggle for the future of the Web is — on the level of theory — a philosophical debate between two opposing theories of knowledge: Heidegger and Carnap. By outlining how each theory has been misinterpreted by its concrete materialization in engineering practice on the Web, we can outline a theory of the future of the Web that goes beyond the current impasse caused by centralization. In other words, from the mistakes of the Web we can outline a philosophy of decentralization. Such a philosophy of decentralization is a much-needed foundational orientation that can prevent future engineering and economic innovation from following the all-too-easy path of centralization. More importantly, it even points at this historical juncture to the role the Web can play in transcending our current era of global political and ecological crisis by renewing the project of the decentralization of knowledge, a project at the heart of philosophy from Socrates to the present day.

The original (decentralized) vision of the Web

Before going into the history of decentralization on the Web, we need to answer the question: What is decentralization? In technical terms, a distributed system is defined by Lamport (1978) as a system with multiple components whose behavior is coordinated by passing messages. Many systems are distributed and in general for a system to be successful there has to be trust between its various components so that if the components are involved in some joint task, each can be trusted to play its role. Examples of technically distributed systems include everything from search engines, where multiple servers work together to find and retrieve data that may be spread out across multiple machines, to the traditional banking system where a single payment on a credit card involves co-operative interactions between the computers of a merchant and a bank. Whereas in distributed systems the components are generally trusted, there is no single trusted authority in a *decentralized system*, and so components have to co-ordinate and negotiate trust separately (Troncoso, *et al.*, in press).

In order to situate the Web, it is useful to take into account the wider context of the centralization of knowledge, although in a vastly simplified form. Distributed systems are not only technical but social. As defined by Hutchins, human social institutions and representations are a form of distributed cognition, where humans share knowledge about the world and themselves via the propagation of representations through various media (Hutchins, 1995). One hallmark of human understanding can then be defined as the use of these representations to guide behavior, including decision-making. For that reason, the enlightenment is defined by Kant as the “use [of] one’s own understanding without another’s guidance” (Kant, 1963). In a centralized system, an authority is in control of another entity, resulting in a loss of autonomy for the controlled entity.

Autonomy can then be defined as the use of one’s own cognitive resources to create and share one’s own representations based on an independent judgment in terms of trust. In some distributed systems, the loss of autonomy may be a reasonable design choice, necessary in order to gain increased powers of co-ordination. After all, one does not want soldiers taking decisions in a battlefield autonomously, or an SQL database deciding on its own what someone’s taxes should be through purely internal random

number generation. Yet as regards humans and their social institutions, centralized control over a fellow human being was seen as biologically natural within the institution of slavery, when bodies were reduced to mere tools in a larger process. However, if one assumes that humans are at least *epistemically equal*, i.e., that all humans have at least the potential to be a member of a community of self-directed knowing subjects (Lynch, 2016), then one can state as the goal of knowledge representation that it should enable humans to strive to be autonomous. If human intelligence is dependent on representations, the ability to navigate and create these representations becomes not just a matter of engineering and education, but of utmost political importance.

A number of justifications of central control have historically been put forward, but until the Enlightenment these were typically based on a claim to some kind of hidden knowledge. To summarize Rushkoff (2010), within Europe this knowledge was generally controlled by the clergy, who monopolized the ability to read and write. With the advent of the Reformation and then the Enlightenment, reading and writing skills spread into the population at large, producing the ability to independently publish and argue over truth and meaning. However, knowledge was still effectively centralized by publishers, who controlled the production of knowledge in the form of books, and the university system (which was one of the few institutions to survive the transition from feudalism into capitalism post-Enlightenment), who controlled knowledge in the form of explicit training and certification. Knowledge itself is a prime reason for control: If someone doesn't know how to do something or how something works, it seems intuitively obvious that they should be put under the control of someone who possesses the knowledge that is proper to the task at hand. Thus, the advent of the Enlightenment led not to a massive decentralization of knowledge but to a re-centralization of knowledge in the hands of a bureaucratic elite, who maintain their power at least in part through their control over knowledge (Rushkoff, 2010). Yet this control could be naturalized, as the time and effort that could be put into the reading and training required to join the "knowledge class" did not seem to scale. To put it crudely, if one wanted access to specific knowledge up until even the 1980s, one would have had to go to Oxford to gain access to the Bodleian library — a task that was simply impossible for the knowledge-starved masses of the earth, who were thus stuck in the proletarian positions of taking orders from the knowledge elite.

After the invention of digital computers in the mid-twentieth century, for the first few decades of their existence these general purpose machines were hidden away like sacred idols by a priesthood of computer operators, with the huddled masses forced to write their programs on punch cards whose answers, in the fashion of a Sibylline oracle, would be given days later. The access to computers by a few was of course aggravating to scientists and a new class of "hackers" who wanted to be able to directly interact with the computer. The breakthrough of time-sharing shattered this monopoly of knowledge (McCarthy, 1962). Time-sharing took advantage of the fact that the computer, despite its centralized single processor, could run multiple programs at once in a nonlinear fashion, making computation much more efficient and accessible. So, instead of idling while waiting for the next program or human interaction, in moments nearly imperceptible to the human eye, it would share its time among multiple humans. Inspired by the spread of time-sharing, the question facing computer scientists was how could computational resources be shared not only throughout time, but throughout space?

The answer, under the auspices of Licklider's tenure at ARPA, was the Internet, and the scientific project to create a "Galactic Network" of researchers that could share computing resources began in earnest (Hafner and Lyon, 1996). After considerable toil, the invention by Cerf and Kahn of a general-purpose protocol for distributed communication, TCP/IP (Transmission Control Protocol/Internet Protocol), led to a plethora of applications that are generally taken for granted today, from e-mail to file sharing. With the military Internet splitting off, the use of the Internet remained from its advent in the late sixties until the late eighties effectively the domain of academic computer science researchers, with little impact on the spread of knowledge outside these rarefied circles.

As the invention of personal computing in the late seventies led to more widespread adoption of computers by the general population, various attempts to turn the Internet into a platform for sharing knowledge began to take shape, with the two most notable being WAIS (Wide Area Information Servers) and Gopher. WAIS was specialized for accessing and searching library indexes, but could be used as a general purpose search engine for searching text on a remote server over TCP/IP. Initially developed by Brewster Kahle, Harry Morris and other programmers at Thinking Machines Inc., WAIS soon became one of the more popular and effective ways to find information on the Internet despite lacking a graphical user interface. Nearly simultaneously, another team of researchers at the University of Minnesota developed another protocol, Gopher, which allowed the organization of information on the Internet through a series of menus that an ordinary person could easily navigate.

Gopher could even be combined with WAIS for effective searching of full text, and it appeared that the Internet was finally poised to create a decentralized digital library of Alexandria. With numbers of users of Gopher and WAIS rising rapidly, the siren song of financial success beckoned. Thinking Machines Inc. stopped allowing WAIS to be used for free, and Brewster Kahle and Harry Morris set up WAIS Inc. to sell the software, which was promptly bought by the commercial Internet service AOL. Likewise, the University of Minnesota decided to start charging licensing fees for the Gopher codebase created by its developers. At the very moment when there was rising interest in the Internet as a potential platform for discovering knowledge by the general public, it seemed as if the first generation of software would put this knowledge behind a paywall.

Luckily, although his paper describing the "World Wide Web" was rejected for the ACM Hypertext conference in December, 1991 in San Antonio, Texas, Tim Berners-Lee decided to go there and give a demonstration. On his way, he stopped at universities and gave demonstrations of how to set up a Web site and "link" using hypertext from one Web site to another. As Gopher and WAIS fell into decline due to the uncertainty around licensing and commercialization, the World Wide Web started to take-off. Both taking key ideas from the concept of hypertext invented by Ted Nelson's Xanadu and earlier systems such as Engelbart's NLS (oNLine System) as well as departing from them, the Web at first seemed rather underwhelming. However, it succeeded because it was both easy-to-use and decentralized.

The first virtue of the Web was a radical simplification of the overly complex academic hypertext systems, allowing broken links and easy-to-use markup in the form of HTML (HyperText Markup Language). Broken links are a fundamental ingredient of the Web which, unlike other existing hypertext systems, does not guarantee access to content. A dreaded 404 error is always possible since no central authority preemptively checks URIs, payloads, continuity of service or even deliver authorization to “mint” them (provided one is in control of a domain name).

The second breakthrough was the layering of HTML hypertext on top of TCP/IP and the domain name system, allowing hypertext “pages” (or rather “resources”) to be identified by URIs (Uniform Resource Identifiers) such as the now familiar <http://example.org>. Berners-Lee viewed this as even more critical to the Web than the use of HTML, since any Web page could link to any other Web page in a decentralized manner and URIs provided a universal space of names so that anyone could buy (or rent) a domain name and create a Web page.

With the easy-to-use language of HTML, the ubiquity of TCP/IP that connected computers all over the globe and the well-understood domain name system for buying names, anyone could easily set-up their own Web site to share knowledge about any subject of their choosing, and thus the Web soon took off as the first truly decentralized system for global knowledge sharing. The Web’s decentralized nature, which allowed anyone to contribute and link to anyone else, made it a “permission-less” platform for knowledge. The decentralized innovation also applied to the core functionality of the Web as developers added new tags, such as the image tag by Netscape, and a constant stream of innovation has characterized the Web ever since its inception. Of course, it helped that CERN was committed to providing the core technology for free and the permission-less innovation was managed by a consensus-run global standards process for HTML, HTTP and URIs at the Internet Engineering Task Force (IETF) and Berners-Lee’s own World Wide Web Consortium (W3C). Still, the Web was not completely decentralized, as the domain name system itself, on which URIs depend, was centralized and requires the licensing of domain names — although once one has bought a single domain name one may host many different Web sites. As regards the decentralization of knowledge, the Web was viewed not as the end, but the beginning: Berners-Lee and others began hoping that eventually it would evolve into a truly universal information space for the sharing of knowledge that went beyond hypertext.

Carnap and the philosophical roots of knowledge representation and the Semantic Web

What would come *after* the Web? Given Berners-Lee’s background as a database administrator at CERN, the obvious next step was to add databases to the Web in a form more amenable to machines than hypertext. The Semantic Web was imagined by Berners-Lee as the next logical step in the development of the World Wide Web, where the Web would go beyond hypertext and connect data in databases. The term “semantic” was used to separate the Semantic Web from the “syntax” of the hypertext Web, including its focus on layout and style that may obscure the knowledge embedded in the Web page. Instead, it was imagined that the data would be unleashed from databases, put on the Web with URIs, and linked together in a decentralized manner.

Berners-Lee’s early thoughts, as given in the first World Wide Web Conference in Geneva in 1994, were that “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” (Berners-Lee, 1994). Having information in “machine-readable forms” requires a knowledge representation (KR [1]) language that has some sort of relatively content-neutral syntax for encoding content (Berners-Lee, 1994).

Under the aegis of the W3C the first knowledge representation language for the Semantic Web, the Resource Description Framework (RDF), was made a W3C recommendation (Hayes, 2004). Interestingly, the first attempt at RDF was thrown out the window by the W3C, who rescinded RDF as a Web standard as it was unclear what the links meant and so how to achieve interoperability (Lassila and Swick, 1999). In the next version of RDF, the Semantic Web was built on a foundation of logical axioms that precisely described the permitted inferences any given statement made. With the help of artificial intelligence researchers such as Pat Hayes and Ian Horrocks, the Semantic Web went from simple links between atoms of data to a full-blown language for knowledge representation. Based on decades of research within artificial intelligence, the Semantic Web was given a formal semantics that could rigorously define any statement in terms of logic, albeit in a tractable manner that kept the language’s expressivity deliberately weaker than first-order logic. It was assumed that the Semantic Web would mature from its original tractable and deliberately weak formulation via standardizing a family of logical languages, from RDF to OWL to first-order logic that could express any and all knowledge fit to be published on the Web.

The philosophical foundations of logic is its own dramatic story, but one that is crucial to understanding both the successes and failures of the Semantic Web. Logic began to develop its new philosophy during the last quarter of the nineteenth century when it was formalized by German logician, mathematician and philosopher Gottlob Frege. Frege had created an artificially restricted language capable of describing the basic operations of logic and tried to use it to provide a logical foundation for mathematics, an endeavor that came to be known as “logicism.” In Ludwig Wittgenstein, one of the most important philosophers of the twentieth century and on many accounts Frege’s heir, we find the claim (both self-defeated, due to the structure of his *Tractatus Logico-Philosophicus*, and hinted at in Frege) that logic should be the privileged medium for *describing the world*. Drawing from both Wittgenstein and the logical atomism of Wittgenstein’s teacher, Bertrand Russell, Rudolf Carnap, arguably the most famous student Frege ever had, took as a starting point these positions (along with the widespread Kantianism found in Germany and Austria) while deeply breaking away from them at the same time.

As one of the founders of the Vienna Circle and the main advocate of the position it advocated known as “logical positivism” or “logical empiricism”, Carnap is perhaps best remembered (unduly so) for his

attempt to eliminate Heidegger's metaphysics via his own approach to philosophy, recast as a "logic of science". Yet, Carnap is a much more prominent philosopher than this crude summary would imply. In his magnum opus, *Logische Aufbau der Welt* (*The Logical Structure of the World*), Carnap (2003) not only described how scientific concepts could be constructed from perception but he also showed how these could be used to methodologically derive concepts. Carnap thus built the equivalent of an abstract machine effecting transformations between inputs and outputs that relied on structural relations rather than content and *relata*.

Contrary to Russell, Carnap's goal was never was to build a metaphysics from sense data or the realm of the given, but rather to show that one could choose an arbitrary basis in one corresponding language from many competing languages and so build up a kind of derivation machine [2] that could be used to derive the relationship of the chosen language to other competing languages. Rather than providing a metaphysical foundation, the "auto psychic" level helped Carnap to bootstrap his system making it liable to be structured in such a way as to allow logical inferences that would connect multiple domains of objects into a coherent logical system. The peculiarity of Carnap's goal was to find an engineering-flavored successor to philosophy.

With Carnap's *Aufbau*, one can see what happens once any notion of content is shed and in its place a system of structural relations is adopted. It is not completely unlike what happened later with the Semantic Web that defined meaning simply in terms of inferential relations, though Carnap's attempt relied on a specific logic at the time and suffered from a host of problems, some of which were identified by Goodman (1977) in his *Structure of Appearance*.

Later, Carnap and the Vienna Circle tried to establish their program on the foundations laid by Wittgenstein in his *Tractatus Logico-Philosophicus* [3]. This endeavor proved less of an exegesis of Wittgenstein's work than an attempt to establish a new philosophical program known as "logical empiricism" by using a framework that, perhaps for the first time in the history of philosophy, seemed to genuinely warrant it.

Traditionally, logic and mathematics had always been a problem for empiricists like David Hume that derived knowledge from perception and the senses. Such foundation, however, did not suit logic and mathematics very well. Wittgenstein took a different approach and introduced the idea that logical statements as chaining of propositions were either true and tautologous or contradictory: In no way did these logical propositions have anything to say about the world. Hence, failure to account for logic and mathematics from experience was no longer a problem since it was no longer necessary by definition. Logic could uncontroversially be incorporated into an empiricist framework without having to either reduce it to experience or to invoke *a priori* transcendental principles as done by Kantian philosophy.

This view of logic leads to a sharp distinction between two kind of propositions: Namely those devoid of meaning, the logical ones, and those that, due to the picture theory of language established in *Tractatus*, receive their meaning from their ability to represent the world by isomorphically corresponding to its basic elements (objects that were concatenated into facts represented in so-called "atomic" sentences). But that was only the starting point that the Vienna Circle needed from Wittgenstein. To fulfill its own goals, it had to somehow "extend" his view of language. As Awodey and Carus (2009) contend, such an extension went in two directions: downwards, to anchor the most basic element of language in sense data, an epistemological task of "interpreting Wittgenstein's 'atomic sentences' as elementary observation sentences"; and upwards, by reaching to mathematics in order to cope with the real epistemological language of science. It was clear for all members of the Circle that accounting for science required escaping the tight boundaries set up by Wittgenstein's propositional logic in *Tractatus*. From the point of view of logic, this was undoubtedly an extension of the tautological character of mathematics, whereas from the point of view of mathematics, this amounted to nothing more than a reduction to empty statements. This two-fold extension was needed for the purpose of articulating the language of the natural sciences in the language of logic, which set the stage for the use of mathematics in the natural sciences thereafter.

Ironically, Carnap's (1959) popular image is that of a philosopher who engaged in a dispute with Heidegger in his famous "*Überwindung der Metaphysik durch logische Analyse der Sprache*" ("The Elimination of Metaphysics Through Logical Analysis of Language"). By contrast, his true persona was much more conciliatory and constructive, taking more after that of an engineer in a standards body than a quarrelsome philosopher. That being said, his critique is interesting for the reflective question it raises. Indeed, his critique of Heidegger relied on the aforementioned distinction between analytic (logical) and synthetic (empirical) propositions, and so Carnap's criteria relegated Heidegger's discourse to the level of meaningless pseudo-propositions since it fitted neither category.

However, applying the same criteria to his own discourse threatened to lead to the same conclusion. "Metaphysical" discourse as well as scientific philosophy potentially endured the same fate. Aware of the issue, Carnap advanced the idea of meta-logic as a means to speak about the propositions of a language in that same language and began to situate his philosophical discourse on the level of metalogic. This move had long-lasting consequences for it meant that after rejecting Hume and Kant and criticizing Heidegger's speech, Carnap now set off to abandon one of Wittgenstein's most important tenets: in his *Logical Syntax of Language*, Carnap (2002) effectively let go of Wittgenstein's prohibition on the use of language to speak about the logical form of language.

Frege and Russell, who both exerted a great deal of influence on Wittgenstein, took the laws of logic to be akin to laws of thought or nature — albeit of a more general status. Wittgenstein, on the other hand, acknowledged that laws of logic were laws of language. Yet, the representational nature of language itself meant that there was only one language representing the world with no possibility whatsoever of stepping outside of it (language itself had become transcendental). To go beyond these limitations, Carnap first had to get rid of the representational view of language (and of meaning at the same time) that was so central to *Tractatus* in the guise of the picture theory of language and trade it for a view that understood language (or rather languages, in the plural) as a calculus based on explicit rules whose

structure can be studied and analyzed by resorting to a hierarchy of other languages: meta-languages, meta-meta-languages, and so on. He also introduced a distinction between the formal and the material modes of speech. The material mode, dealing with objects and facts, suited science inasmuch as it had an empirical import. Yet, the elucidation of both logic and the whole of knowledge, the task of philosophy according to Carnap at the time, required the formal mode. Philosophy reframed as the (meta)logic of science could at long last find a proper place to set in.

The unity of language was henceforth broken as the rules of syntax no longer obeyed any univocal representational imperative and could be engineered at will to fit the needs of a formal articulation of the content of science. The “principle of tolerance”, exposed in *Syntax*, stated that “everyone is at liberty to build his own logic, *i.e.*, his own form of language, as he wishes. All that is required of him is that, if he wishes to discuss it, he must state his methods clearly, and give syntactical rules instead of philosophical arguments.” (Carnap, 2002) It became (and remained) the core principle behind Carnap’s enterprise to his death [4].

One year after the original publication of *Syntax* in German, Carnap, impressed by Tarski’s definition of truth, decided to follow him and adopt the so-called semantic approach. This is far too convoluted a story to tell in its entirety, but a few words will suffice. As we’ve seen in *Syntax* Carnap had restricted his meta-language to purely syntactic terms. Tarski’s definition of truth showed that it was possible to propose a sound characterization of truth that used descriptive terms. The elimination of the picture theory of language based on atomic sentences had entailed the elimination of meaning in *Syntax* in favor of syntax and the formal mode of speech of the logic of science. Still, Carnap soon recognized that “the restriction to the syntactic method was just inappropriate for the logical analysis of science” (Wagner, 2015) [5]. Carnap’s main originality here, and his immediate legacy for AI, was precisely his motivation for adopting (up to a certain point) Tarski’s semantic approach:

“the difference between Tarski’s method and my method of semantics is to a large extent to be explained by the fact that Tarski deals chiefly with languages for logic and mathematics, thus languages without descriptive constants, while I regard it as an essential task for semantics to develop a method applicable to languages of empirical science. I believe that a semantics for languages of this kind must give an explication for the distinction between logical and descriptive signs and that between logical and factual truth, because it seems to me that without these distinctions a satisfactory methodological analysis of science is not possible.” [6]

In contrast to Tarski’s avowed aim to shed light on the methodology of deductive sciences, Carnap’s interest was elicited by the possibility of extending this newfound framework to all sciences, including the empirical ones, on the condition that a criterion separating logical truths from factual ones may be found — or rather *devised* — for all languages under consideration.

Semantics continued to evolve after Carnap. The standard approach to formal semantics is now commonly known as “model-theory” or “Tarskian semantics.” For a while though, Carnap’s semantics differed noticeably from model-theory and did not follow all of Tarski’s technical developments. When artificial intelligence (AI) researchers turned to model-theory they did it in a spirit which was much more reminiscent of Carnap than it was of Tarski. After all, their systems purportedly dealt with “descriptive” terms — actually, that was their main benefit [7]. What KR and the Semantic Web call ontologies is a way to deal with such descriptive terms in a logical fashion, by treating them as pre-defined axioms or “meaning postulates” (Carnap, 1952). In that vein, Nicola Guarino, a scholar known for having established bridges between philosophy and ontology engineering, defines an ontology in the following way:

“to specify a conceptualization is to fix a language we want to use to talk of it, and to constrain the interpretations of such a language in an intensional way, by means of suitable axioms (called meaning postulates). For example, we can write simple axioms stating that reports-to is asymmetric and intransitive, while cooperates-with is symmetric, irreflexive, and intransitive. In short, an ontology is just a set of such axioms.” (Guarino, *et al.*, 2009)

By building on top of these axioms the goal of the nascent field of artificial intelligence in its early days was to build real-world applications to interact with their environment and take concrete action therein based on logical inference. Such a move is not alien to Carnap’s philosophy. Quite the contrary. In parallel to his studies in semantics in the 1940s, Carnap did advance a notion of “explication” which amounts to substituting for a term to be explicated (the *explicandum*) another term, “given by explicit rules for its use” (the *explicatum*). Of paramount importance is the fact that “the *explicandum* may belong to everyday language or to a previous stage in the development of scientific language” [8], since that allowed to extend the formalization outside the boundaries of science so as to include common sense as well. With the addition of that small caveat (and considering that French epistemology per Foucault has shown that there is no “precursor” whatsoever in the history of science!) Carnap really stands out as the forefather of AI and knowledge engineering.

Knowledge engineering is a branch of AI that, building on ideas elaborated by Carnap, took its inspiration from philosophical ontology and certain strands of metaphysics (especially the Aristotelian school, Husserlian mereology, and analytic metaphysics). Philosophically-inspired knowledge engineering (which does not account for the whole discipline) has long espoused realist views in metaphysics to fulfill its own need to formalize axioms with higher-order principles, thereby setting a hierarchy between formal descriptions of domains and top-level ontologies inspired by previous philosophical work. Top-level ontologies manifest themselves as unifying principles of modeling whose overarching conceptualizations are buttressed by the practical need of “interoperability” between various formalizations of domains such as biomedicine [9]. Yet the field of knowledge engineering was held back

by a lack of agreement regarding the “correct” upper-level ontology, as should be expected given the difficulty inherent in combining the development of standards by committee with the well-established lack of consensus in metaphysics.

With the advent of the Semantic Web, the principles of decentralization could be applied to knowledge engineering. Devised by AI researchers directly influenced by Carnap, the Semantic Web adopted Carnap’s tolerant viewpoint that different logical languages could co-exist. The Semantic Web avoided any need to choose a single top-level ontology, but instead allowed anyone to create an ontology and post it to the Web simply by associating the terms of these “vocabularies” with a Web proper name; in other words, a URI. These ontologies were capable of logical inference powered by a Tarski-style formal semantics outlined in the W3C standards for RDF, although subsequent logics developed by the W3C given by a host of standards around the Web Ontology Language (OWL) fractured the Semantic Web into various “stacks” of knowledge representation languages, each with its own different formal semantics. The vision was that the Semantic Web would decentralize knowledge engineering, allowing data from everything from spreadsheets to databases to seamlessly connect on the Web via formal ontologies that would organically grow over time.

From Heidegger to knowledge graphs and machine learning

Despite its promising decentralized vision and strong foundations in knowledge engineering, the Semantic Web effort stalled in terms of practical uptake. While the hypertext Web had within a few years produced an exponential growth in Web sites, the Semantic Web mostly produced an exponential growth in academic papers with little real-world impact. Frustrated, Semantic Web stalwart and original co-designer of HTML with Berners-Lee, Dan Connolly (2006), made a fascinating observation: The “infoboxes” of Wikipedia contained data that was both inherently stable and crowd-sourced. While the professional ontologists working on the Semantic Web failed to produce real-world schema usage outside a few limited domains such as biomedicine, Wikipedia offered a ‘crowd-sourced’ schema for almost all of reality that was updated by volunteers at zero cost. Seeing the opportunity to turn Wikipedia into a structured database of knowledge that could power new AI applications, a startup called Freebase formed to “scrape” Wikipedia’s infoboxes and create a dynamically updated and curated knowledge base. Simultaneously, a number of German Ph.D. students working on the Semantic Web started creating DBpedia, by converting Wikipedia’s infoboxes to the Semantic Web language RDF (Auer, *et al.*, 2007). While Freebase kept their curated version closed, DBpedia was an open Semantic Web database. As it was available to the research community and smaller companies, DBpedia inspired a wave of revived research on knowledge engineering. However, what was less noticed was that Google quietly acquired Freebase, and then soon began hiring experts in knowledge engineering, including R.V. Guha, one of the key designers of RDF at the W3C and pioneer in applying knowledge engineering to artificial intelligence.

Another way to harvest knowledge was to have users manually add structured data to existing Web pages, with Wikipedia ‘info-boxes’ being just one example. For example, a store’s Web page could explicitly represent its opening hours and location as some kind of formal, machine-readable knowledge. Microformats, started by Tantek Çelik, made it easy for users to add data such as their contact information and calendar in a structured way to their own Web pages. The W3C standard Gleaning Resource Descriptions from Dialects of Languages (GRDDL) could convert this to Semantic Web formats, and soon enough the W3C developed a competing standard for embedding Semantic Web metadata into Web pages called RDFa, although its success was more limited due to the lack of adoption of metadata by developers (Connolly, 2006).

However, Yahoo!’s engine started indexing and processing this structured data on the Web in order to customize search results, such as showing opening hours of stores and phone numbers on the main Yahoo! search-page rather than requiring a user to ‘click’ on a link to get this valuable information (Mika, 2008). Other search engines soon followed, including Google with Google Rich Snippets. This led to an explosion of structured data on the Web, as Webmasters thought that adding structured data would help their search engine optimization. The editor of HTML5, Ian Hixie, created yet another incompatible general purpose standard for embedding data called microdata. Although a “lower-case semantic web” was taking off in the form of semantic annotations to existing Web pages, it seemed competing standards and search engines were causing massive fragmentation of this kind of “low budget” data embedded in Web pages.

In 2011, Google’s plans for the Semantic Web became clear: Led by Guha, Google had created a massive knowledge representation framework called “schema.org” to unite the fragmented structured data present on the Web (Guha, *et al.*, 2016). Using the considerable clout of Google, other search engines such as Yahoo!, Microsoft and Yandex joined the effort so that every search engine could consume the same kinds of structured data, and Web page authors would know which logical terms to use in order to add knowledge to a Web page. Although not an open standard and controlled informally by a small group of search engines, schema.org finally made structured data take off, so that soon up to 10 percent of the Web was using structured data. Even the Facebook “Like” button began embedding data using RDFa, making structured data ubiquitous on the Web.

For the most part ignoring top-level ontologies based on metaphysical distinctions and even formal semantics, Google designed various lower-level ontologies for domains of interest to search engines, such as e-commerce, movie, and music information. A social process based on mailing list discussions was put in place to add new ontologies, and soon schema.org was growing to encompass more and more of the world’s knowledge. Thus ironically, much of the academic work on logical inference and formal semantics thought to be needed by the decentralized Semantic Web ended up being ignored, while a human-powered yet centralized web of knowledge began taking off. Furthermore, Google was

using schema.org and Freebase's version of Wikipedia to create their own internal version of the Semantic Web (or more precisely, of the Linked Data Cloud), called the Google Knowledge Graph. This proprietary database was put in place to connect the vast variety of heterogeneous knowledge spread throughout Google's various online services. At the same time, other companies such as Yahoo!, Microsoft and even Apple started creating their own competing proprietary knowledge graphs. The use of these knowledge graphs started becoming increasingly common in new products. Behind Apple's Siri's knowledge of the world lies the formal knowledge engineering of a spin-off of Stanford Research Institute (SRI) that formed the foundation for Apple's knowledge graph.

One of the most long-standing problems of knowledge engineering was how to dynamically add new knowledge. Although schema.org and Wikipedia solved this by having humans essentially crowd-source knowledge, the deluge of data released by the Web far outweighed the cognitive resources of even crowd-sourcing. After all, it seemed infeasible to pay people to identify those in every single photo on Facebook, and users identified people explicitly in a minority of photos. Also, as reality changed so did knowledge itself, and there were simply not enough knowledge engineers to manually update various knowledge graphs to take into account every single change.

The answer was obvious: Computers had to be able to learn knowledge themselves, both with and without human supervision. Also coming out of artificial intelligence, the field of machine learning had been developing quietly in parallel to the more traditional knowledge engineering approaches. Machine learning, while it had some early successes in the work of AI pioneers such as Selfridge, had always suffered from a lack of data ... until the Web came along. Machine learning itself found in the form of the Web a massive input data set that was increasingly updated in real-time.

Although the techniques behind machine learning seemed deceptively simple, by virtue of having as input data a massive representation of collective human existence, these simple techniques could tackle problems beyond knowledge engineering, with machine translation as the example par excellence. While techniques based on knowledge engineering and logic produced terrible results, when such word-by-word translation based on the senses ("semantics") of words were replaced by "phrase-based" statistical translation that took advantage of parallel corpora, projects like Google Translate could soon produce passable translations of many languages. The same general approach of relying on real human data rather than formal rules and logical inference also applied to varied fields from speech recognition to search engines. As reportedly put by Frederick Jelinek, "Every time I fire a linguist, the performance of the speech recognizer goes up." Although knowledge engineers were experts in transforming human knowledge into formal representations, machine learning experts would rather throw explicit human knowledge out the window and look for the knowledge implicit in the data itself.

As the adage in machine learning circles goes, "There's no data like more data." Yet storing and processing data did not come without costs. With the amount of data on the Web skyrocketing into the millions of terabytes, what ended up mattering for the future of the Web was the ability to handle data that was larger than could be fit on a single machine, which in turn required large distributed — but centralized — data centers to handle. In other words: "big data." The machine learning field blossomed, producing astonishing success powered by tweaks to a relatively small number of algorithms ranging from the simple Naive Bayes that calculates the probability of data fitting a pre-set number of classifications, to more subtle support vector machines that could project data to a higher dimension. As the ability to handle "big data" and to understand these algorithms was outside the capabilities of many users, the center of power on the Web moved to those that had the data centers and machine learning expertise. While the Semantic Web imagined a vast web of knowledge representations structuring the world's data, what had actually ended up happening was that a few companies had literally developed copies of the entire world's data, and by cleverly applying algorithms to this unstructured data, they were able to extract immense amounts of both knowledge and wealth by predicting patterns in everything from user buying habits to results of elections. The foremost company in this space was Google, a self-declared AI company. However, the kind of AI championed by Google had a far different philosophical heritage than Carnap's idea to formalize all human knowledge. Strangely enough, Google was the child of a philosophical heresy inside AI, a heresy based on Carnap's nemesis Heidegger.

Of course it was not so strange, as nearby Berkeley was the home of Hubert Dreyfus, a philosopher who had brutally critiqued AI (and implicitly, Carnap) by relying on a Heideggerian critique of logic in his *What Computers can't do: A critique of artificial reason* (Dreyfus, 1972). In his book, he claimed that AI was impossible as human intelligence required the full process of growing up in a human body. This book was not only a philosophical riposte against using knowledge engineering techniques to create a human-level artificial intelligence, but the book was based on a RAND report, *Alchemy and artificial intelligence* (Dreyfus, 1962). This report — and others like the Lighthill one — was influential in determining if AI should continue to receive massive national government funding of the kind the Internet was receiving. The answer of Dreyfus was a resounding negative. Combined with official reports such as the Lighthill report, the impact of his Heideggerian critique almost caused AI funding itself to be halted, during what was dubbed "AI winter" in the 1970s to early 1990s. Therefore it should not be surprising that the reading of Division I of *Being and Time* given by Dreyfus defined not only a generation of philosophers but also AI researchers, given that Heidegger helped destroy their research funding.

While some AI researchers reacted very negatively and repeated Carnap's attack that Heideggerian philosophy was nonsense, others were quietly absorbing the insights of Heidegger and incorporating them into a new kind of AI research (Dreyfus, 2007). Of the number of clearly Heideggerian critiques Dreyfus makes of artificial intelligence, the one that was most thoroughly taken to heart by many in the artificial intelligence community was the non-representational nature of knowledge and problem-solving. Early work attempted to replace explicit logical frameworks for representing knowledge with numerical computation over connected graphs of nodes. In these graphs, the nodes were labeled "neurons" and the whole line of research branded "connectionism" insofar as it tried to stay as close to the neural as possible but slowly evolved into a general learning framework (Smolensky, 1988). Staying close to the neural level was a difficult task, given that if anything neuroscience shows that the electrochemical process in the brain is quite removed from number crunching, so eventually AI researchers simply gave

up on modeling neurons and generalized neural networks into more generic machine learning algorithms, based on anything from pure ad hoc design to a more principled Bayesian framework. In a sense, machine learning as a field is a strange offspring of the influence of Heidegger (Dreyfus, 2007).

Against Carnap, Heidegger believes the idea of knowledge as logic and facts comes only via the “theoretical attitude,” a mode of detachment that parasites on and often misinterprets this lower level of deeply involved engagement in the world. The overwhelming importance of the practical task in defining the very world we live in was taken up in a Heideggerian context not only by Winograd, but by the forefathers of ubiquitous computing such as Mark Weiser, and so slowly but surely became second nature to computer engineering (Weiser, 1991). The eminently practical part of Division I of *Being and Time* cuts even deeper than Dreyfus realized. One possible, if simple-minded, rejoinder to Dreyfus’ contention that intelligence requires embodiment would simply be to build a *body* in the form of a robot (Brooks, 1991). Still, Heidegger was clear that what matters about embodiment was not the sheer presupposition of having a physical body that can interact with the world, but that embodiment enables *having a world*, and it is this *worldhood* (*Welttheit*) that is defining of being (Heidegger, 1962). A *Dasein* is defined by its intentionality and the various practical projects that it engages in, from hammering a nail to posting a photo online, and these activities transform the entities into equipment (*Zeug*) that can then be used to accomplish tasks, and so the entities are “ready-to-hand” (*zuhanden*) (Heidegger, 1962). When an object is “ready to hand”, it becomes the exact opposite of explicit knowledge with various attributes that must be consciously grasped: More like a hammer in use by a skilled carpenter or a screen when a programmer is debugging, the tool itself becomes invisible by virtue of being thoroughly integrated into the practical activity itself. It was this insight that ended up transforming AI and machine learning more than any other insight from Heidegger, even if the insight was perhaps altered beyond recognition to Heidegger himself in transmission.

Few researchers had made the “Heideggerian revolution” in AI explicit until Terry Winograd, with the help of Chilean economist and exile Fernando Flores (1986), authored *Understanding computers and cognition: A new foundation for design*. Under the influence of the early Heideggerian metaphysics from *Being and Time* as channeled into the Anglophone world by Dreyfus, as well the strange idiosyncratic cybernetics of autopoiesis from Maturana, Winograd and Flores attacked the logical foundations of artificial intelligence (one of which we saw comes from Carnap) and explicitly gave a new metaphysics for computing. Reinterpreting Heidegger and Merleau-Ponty in terms of a distinctly computational framing, Winograd and Flores began integrating the human into the heart of the computational system itself.

Rather than attempting to create a third-person scientific perspective or an autonomous artificial intelligence based on logic, Winograd and Flores turned to a metaphysics of human-centered design, where the central task was transformed from representing human knowledge to using a machine to better enable the implicit and embodied knowledge of humans. Thus, logical representations of the sort championed by Carnap and AI knowledge engineers were considered *passé* and metaphysically suspect at best. The key to the new Heideggerian metaphysical foundations for design was Heidegger’s concept of *Zuhandenheit*, of “ready-to-hand”, where the goal was to have the computational apparatus become completely transparent — invisible — to the human. If there was to be some kind of technical breakdown, the technological apparatus was to become reshaped based on human feedback with the ultimate goal of re-establishing its own self-organization that continuously improved in the face of the messiness of the world. In Heidegger as well as the theory of autopoiesis, there was no meaning outside of the phenomenological world, and so formal semantics and the rest of the Carnap-inspired apparatus was thrown out, with a new emphasis being placed on learning and human-computer interaction.

Machine learning itself was only fleetingly approved by Winograd and Flores (1986), as they noted that “there is new interest in the phenomena of learning” due to the fact that “formal analysis” was “too limited to form the basis for a broad theory.” Unlike Dreyfus, Winograd and Flores felt that AI should not be built on neural theories, as “detailed theories of neurological mechanism will not be the basis for answering the generation questions about intelligence ... any more than detailed theories of transistor electronics would aid in the understanding of computer software.” In their radical re-interpretation that blended together Heidegger and cybernetics, the human was to become part of a new kind of distributed cognitive system that continually learned from its mistakes. Technology aimed for ever smoother, and eventually invisible, integration with the human. In other words, Winograd and Flores had laid the metaphysical foundations for Google.

Perhaps, then, Carnap’s project to make knowledge explicit through “explication” is a possible corrective to the invisible and implicit role of Heideggerian-inspired AI — and so the realization of Carnap’s influence on AI may eventually turn out all the more intriguing than his critique of metaphysics. Both Carnap and Heidegger were against declaring one “top-level” ontology or master meta-logic. The dispute between Carnap and Heidegger overshadows the positions they unexpectedly share about the importance of practical efficacy in determining the proper framework for “representing” knowledge. In fact, Heidegger and Carnap had met and had cordial discussions in the 1920s and Carnap had carefully read *Being and Time* (Friedman, 2000). Despite the language employed, Carnap even shares some commonality with Heidegger’s own critique of traditional metaphysics [10].

In philosophical circles, Carnap’s approach to ontology is widely characterized as “deflationist”. In “Empiricism, semantics and ontology” (Carnap, 1950), Carnap dissociates what he calls “external questions” from “internal” ones. The latter deal with what exists in a given linguistic framework (numbers in mathematics, atoms in physics, etc.) while the former deal with existence *simpliciter*, questioning the framework itself. For Carnap — and this is reminiscent of Wittgenstein’s insight — one cannot step outside of all the available linguistic frameworks and meaningfully articulate such meta-questions. On the other hand, Carnap still contends that an evaluation of the frameworks themselves is possible, but only a practical one because no theoretical evaluation of a linguistic framework is available within the same linguistic framework. Although, properly speaking, the evaluation may be both theoretical and practical, for weighing the consequences of formal apparatuses may be part and parcel of the contribution of other disciplines, each defining a different linguistic framework. As regards

computer ontologies, those disciplines include HCI for instance, to which Winograd himself contributed after his turn to Heidegger.

The pragmatic element of Carnap's thought, which he did not explore further himself, nevertheless plays a central role in his philosophy. The difference between the two remain that Carnap's evaluative scheme is still framed in scientific terms whereas it is rooted in "human" experience via *Dasein* in Heidegger. Carnap (1937) may have forfeited the world to the explore "the boundless ocean of unlimited possibilities" [11], disclosed by his principle of tolerance, yet his pragmatism re-anchors his heritage on more worldly ground, possibly easing the discussion between the two rival schools of philosophy at long last. When AI researchers turned to HCI and machine learning after reading Heidegger, they betrayed a tendency which could have stemmed from Carnap himself due to his focus on practical efficacy!

If anything, the Knowledge Graph prompted machine learning towards the "deep learning" revolution. The problem was how to connect the unstructured data and classification tasks that machine learning excelled at with the kinds of complex structures embedded in knowledge representations. Taking image recognition as a paradigmatic example, an image contains not only figures, but also these figures contain faces, which in turn contain eyes and mouths. Or in the case of speech, it was from recognizing elementary phonemes that a machine learner could build entire words, then named entities, and then phrases, sentences, and paragraphs — and finally to place the text in some library-like hierarchy of subjects. These kinds of features that involved multiple and hierarchical features were at first impossible for machine learning. However, due to the work of pioneering AI researchers like Geoffrey Hinton (now at Google), layers of neural networks were hooked together, where each layer could recognize specific features and guide the learning of not only itself, but other layers via feedback (LeCun, *et al.*, 2015). These cascades of machine learning algorithms became known as deep learning algorithms due to their ability to learn at many different levels of abstraction simultaneously. Although computationally even more expensive, these deep learning algorithms formed the magic glue that could connect unstructured data, such as photos, videos and text in books and Web pages, to the structured Wikipedia-style knowledge of the "real world" stored in each company's knowledge graph. The numbers of inputs to these machine learners started scaling to billions, far more than what could be handled without a data center.

As an aside, it must be noted that Carnap devoted the most important part of his career (from the 1940s to 1970) to elaborating an inductive logic, a project that remains quite obscure especially when compared to his previous efforts. Yet, a parallel has sometimes been drawn between Carnap's inductive logic and the reinforcement algorithms used in machine learning (Kreinovich, 1993, 1992). Scholarship on Carnap hasn't caught up with such insights but one can only hope it does in the foreseeable future.

Heidegger himself would have recognized a strangely familiar and monstrous return of his philosophical enemy, enframing (*Gestell*), in the knowledge graph. The knowledge graph conceives of the world as facts attached to "objects" that are always "present" due to the Internet and so framed due to the ubiquity of knowledge graphs and machine learning. In Heideggerian terms, knowledge graphs are a formalism to represent not the properly ontological, but the merely *ontic* — the world as facts. This attempt to define the world as entities with properties and concepts to be calculated over by machine learning algorithms, with Being somehow being at the top of the hierarchy, is for Heidegger an *ontotheology* par excellence that attempts to enframe a particular conception of the world as historically eternal, and so squarely violates the metaphysical stance of his later years. If he were alive today, Heidegger would no doubt point out that by virtue of regarding the world as a collection of entities with definite and objective properties, the knowledge graph shows itself to be wedded to the classical Platonic and Cartesian traditions that ignore the question of the "meaning of Being," a question that can only be answered by a *Dasein* that is "thrown" into the world (Heidegger, 1962). Regardless of this misreading, a bizarre if unrecognized neo-Heideggerian ambiance pervades Silicon Valley, from the emphasis on user experience to the disappearance of the interface into an array of sensors that are directly placed on the body. Due to this invisibility, the mobile phone, and so Google, becomes a literal extension of our own knowledge, and it becomes unclear how we would even function without it (Halpin, 2013). Only when the phone is absent or malfunctioning do we notice how utterly dependent we have become on the Internet for our knowledge. Thus, both the technical efficacy and the political problem given by Google's post-Heideggerian philosophy become apparent.

Conclusion

Where does the rise of deep-learning and proprietary knowledge graphs leave us in terms of decentralization and autonomy? It seems that knowledge on the Web is now more centralized than before in a few large providers such as Google and Apple, even if the human capabilities enabled by data-driven applications such as Google Maps and Siri are far more superior to those provided in the "golden days" of Web pages and Wikipedia. As a few companies currently control the massive amounts of computing power, closed algorithms and massive data sets needed to fuel the machine learning algorithms that operate behind the scenes in these new applications like Google Maps and Siri, so likewise only a small elite can truly harness the potential power of data on the Web. There is now widespread concern that this vast power may be abused, and there is spreading among the general population a fear of these companies and a distrust of the Internet (Morozov, 2014). Can the Web return to being a tool of empowerment?

In the era of Diderot's *Encyclopédie*, knowledge was bound to the function of every tool: An axe for cutting, looms for weaving and even dyes for wig making all featured prominently in the *Encyclopédie*. In the transition to the Web as a universal space of information, the truly necessary tool is the universal abstract machine, the Turing Machine that executes any computable algorithm. As through education and literacy our ancestors learned how to autonomously extend their physical capabilities with modern tools

and learned how to autonomously organize in a larger complex social fabric than simple face-to-face meetings, through programming humans can learn how to communicate with the machines necessary to autonomously understand and control the complex technological world we have inherited. In this regard, programming is not simply the learning of a particular programming language, from Lisp to HTML and JavaScript. What is necessary is for the generalized skillset of scientific, logical and algorithmic thinking that underlies programming to be spread throughout the population. This does not mean it should in any way supersede our previous languages and modes of thinking, just as writing did not absorb non-verbal tool-use and the visual arts, but that it is necessary in order to maintain autonomy in the era of the Internet. Rather than a valence of description of the world or a technique for controlling the world, it would be far better to think of algorithmic thinking as yet another capacity that can be developed and nurtured in future generations due to its own limited yet powerful capacity: A meta-language for controlling the general abstract machines — computers — that currently form the emerging global infrastructure of much of our inhabitation of the planet. Without the ability and freedom to navigate through these programs, autonomy would be lost.

Tim Berners-Lee has, via the Semantic Web, long advocated for open data. It is now obvious that open data is necessary but not sufficient for the development of autonomy. The ability to think algorithmically and program is useless in terms of the decentralization of knowledge unless the proper infrastructure is provided. Decentralized open data and even open versions of the knowledge graph like DBpedia are rather meaningless in terms of human knowledge if only a small minority controls the data centers and machine learning algorithms needed in order make sense of the data. The Semantic Web should encompass more than just open data! If the key to the autonomy of future generations is control over knowledge, then not only must there be open access to the data such as provided by Wikipedia and DBpedia, but there must also be control of the data centers and algorithms by ordinary people. Data centers are already becoming increasingly cheap to deploy due to the Cloud, but they are still fundamentally controlled by corporations rather than people. Likewise, the machine learning algorithms that appear currently as radically opaque trade secrets need to become open algorithms that can be inspected and deployed by anyone. Decentralization must mean seizing back control not only of data, but of algorithms and data centers from centralization, which is a political task for the future. This political task has been latent in philosophy for decades, as articulated by Lyotard (1979) in *The postmodern condition*: "Give the public free access to the memory and data banks."

Autonomy and decentralization does not involve "de-linking" our individual lives from Silicon Valley, just as there was no "de-linking" of the "Global South" from the global market. For better or worse, our immediate survival is evermore tightly knit to the infrastructure which truly has a second nature. It makes little to no sense now to simply let go of it despite the fact that it has become avowedly unsustainable and its functioning increases the strain on the planet, a decaying infrastructure whose current energy-intensive trajectory leaves it to endure the same fate as more ancient dwellers of the biosphere: Extinction. What we do with other beings which face a common threat, those which contribute to futuring as well as those who contribute to defuturing (Fry, 2009), remains to be seen; however, we can imagine that data centers under popular control can be decentralized, and ultimately made more ecologically sound. It should be remembered that the choice between letting go of or embracing digital technology may not even be a fair one. For, if predictions are true, lack of affordable access to oil and resources will drastically reduce the possibility of maintaining, repairing and adding new infrastructures within just a few decades.

The choice between Silicon Valley gadget-making (which amounts to planned obsolescence on a stratigraphic scale since new geological strata currently produced are made out of the plastic and exotic metals excreted by discarded digital devices) and radical technological abstinence is misguided. Instead, to borrow an expression from Donna Haraway, we might have to learn to "stay with the trouble" meaning that not only ourselves but our world has been hybridized with technology (Haraway, 2016). Just as we may not have the choice of leaving our decaying planet to settle on Mars or any other fanciful destination desired by Elon Musk, neither may we have any other choice than to understand and seize control of the knowledge graphs and algorithms which are part and parcel of the global infrastructure. Global — centralized — answers provided by oligopolies follow patterns which do not disclose adequate conditions for the necessary reorganization required if we are to collectively survive the coming hardships, notwithstanding the heroic posturing of Elon Musk: After all, Musk's master plan for 3,000 years is simply to escape the planet and extract whatever is needed to achieve this goal beforehand! To escape the practical problems posed of the present due to the possibility of a "solar catastrophe," as posited by Lyotard in his later years, is escapism. Instead, the intellectual reserves of philosophy should be aimed directly at the problems of political and ecological sustainability inherent in our infrastructure, including the infrastructure of knowledge given by the Web (Lyotard, 1988). For example, given the global scope of climate change and the need for better scientific data collection on carbon emissions by individuals, it is more likely that a decentralized Web in the hands of an empowered population will be crucial for the future.

The future of the Web will then involve a radical practical re-design of machine learning, data centers and knowledge representation in order for knowledge to become truly decentralized. The vision of the Semantic Web should not be caught up in arguments over logical frameworks, but focus on the elements of what it would take to empower people with knowledge: Not just data, but kinds of thinking and infrastructure. The future will then be more contentious than just opening up datasets. The decentralization of knowledge is a political struggle for power over knowledge in the context of an ecological crisis, it is a re-appropriation of what Tony Fry calls "future-making" so as to multiply the ongoing experiments out of which answers (in the plural) may eventually emerge and scale. That is the crux of decentralized knowledge: It must fit local conditions and globally scale whenever needed at the same time. In order to rescue the potential of these technologies, we should rescue their potentials given by philosophy. Both the Semantic Web and Carnap foresaw a future where all of the world's knowledge could be self-organized without a "master plan" but still ultimately strive to be communicated. Carnap's tolerant or decentralized vision of multiple — and possibly incommensurable — languages being developed to aid in large-scale distributed cognition can be extended outside of the confines of the logic

and model theory of the traditional Semantic Web, in order to encompass the opening and sharing of machine learning algorithms, thus providing new frontiers for knowledge engineering itself.

Although the struggle to re-decentralize the Web will have both a technical and political dimension, the heart of Carnap's work could be considered political, and Carnap had hoped that the spread of knowledge would eventually eliminate war, poverty and obscurantism. Likewise, Heidegger's insights into the powers of phenomenology and the primacy of non-conceptual content, as well as his warnings over the dangers of having logic and calculation rule over life, can be heeded without falling into the trap that led him to consciously support Nazism. As we have shown, these dueling philosophical traditions have had a remarkable, if underground, influence on the current centralized Web. These philosophical insights need to be updated and consciously synthesized into a philosophy of decentralization, a philosophy of the Web that can ground the political and technical tasks of the twenty-first century. It's now or never.

About the authors

Harry Halpin is a researcher at Inria (*Institut national de recherche en informatique et en automatique*) currently coordinating the NEXTLEAP project to investigate next-generation standards for security and privacy. Previously, he worked at the World Wide Web Consortium (W3C) on standardization around cryptography and the Semantic Web. He received his Ph.D. from the University of Edinburgh in Informatics under the supervision of the philosopher Andy Clark.
E-mail: [harry \[dot\] halpin \[at\] inria \[dot\] fr](mailto:harry@halpin.inria.fr)

Alexandre Monnin is a researcher at Inria Sophia Antipolis (UCA, Inria, I3S, WIMMICS). A philosopher, Alexandre Monnin received his Ph.D. from Paris 1 Panthéon-Sorbonne University. He has been working on the architecture and philosophy of the Web. In 2011, he initiated the French DBpedia project. Since 2013 he is a member the network of experts of Etalab, the Prime Minister's open data agency. He is also the architect of Lafayette Anticipation's (the Galleries Lafayette contemporary art foundation) digital platform.
E-mail: [alexandre \[dot\] monnin \[at\] inria \[dot\] fr](mailto:alexandre@monnin.inria.fr)

Acknowledgments

Harry Halpin was supported by NEXTLEAP (EU H2020 ref: 688722). The authors would like to thank Yorick Wilks for his thorough review and Francesca Musiani for her patience.

Notes

1. "Knowledge representation" is used as a technical term in artificial intelligence for the formalization of knowledge into computational frameworks.

2. Moulines, 1991.

3. *Tractatus* (Wittgenstein, 1922) is available online, in the original German version side by side with two canonical English translations: <http://people.umass.edu/klement/tlp/>.

4. See in particular (Creath, 2009) and (Monnin, 2015) for a discussion of the relevance of this principle to understand KR and the Semantic Web.

5. As explained by Carnap himself in his "intellectual autobiography", in Schilpp (1963), pp. 59–60.

6. *Ibid.*, p. 932, quoted in Wagner (2015).

7. See in particular Hayes (1995) and McDermott (1978). The semantics of RDF, the basic building block of the Semantic Web devised by Hayes, is based on the so-called "Tarskian semantics" (Hayes, 2004), which was revised in Hayes and Patel-Schneider (2014). A more apt designation would be "Carnapian semantics + models"! Very recently, AI pioneer Patrick Hayes has come to the same conclusion in an exchange with philosophers and engineers on a specialized mailing list: "It is true that Tarski did not say much about how to use formal languages to describe the real world. But Carnap certainly did. '*Topologie der Raum-Zeit-Welt*' (1924) for example. Just take a look at the later parts of his '*Introduction to Symbolic Logic and its Applications*' (1958) which presages almost all of what is now called formal ontology design. (It was reading the last one as an undergraduate in 1965 that pushed me to get into AI, by the way. Carnap did a better job of logical KR than McCarthy ever did.)"; see <https://groups.google.com/forum/#!msg/ontolog-forum/ux1EFGQdfNE/JLFmxkoPAqAJ>.

8. The most complete exposition of Carnap's concept of explication can be found in Carnap (1962). For a reading of Carnap's philosophy through the lenses of the concept of explication, see Carus (2007). See also Wagner (2012) for further discussion. Monnin (2015) stresses the importance of explication to link Carnap to the logical AI movement and more recently, to the Semantic Web.

9. Guarino, *et al.* (2009): "For practical usage of ontologies, it turned out very quickly that without at least (...) minimal shared ontological commitment from ontology stakeholders, the benefits of having an ontology are limited. The reason is that an ontology formally specifies a domain structure under the limitation that its stakeholder understand the primitive terms in the appropriate way. In other words, the ontology may turn out useless if it is used in a way that runs counter to the shared ontological commitment. In conclusion, any ontology will always be less complete and less formal than it would be desirable in theory. This is why it is important, for those ontologies intended to support large-scale

interoperability, to be well-founded, in the sense that the basic primitives they are built on are sufficiently well-chosen and axiomatized to be generally understood."

10. Gabriel (2012, 2009).

11. On the formula and Carnap's philosophy being "world poor," to adopt an expression from Heidegger, see Mormann (2010) and Monnin (2015). The pragmatic element of Carnap's thought is given due consideration in Uebel (2012).

References

Sören Auer, Christian Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak and Zachary Ives, 2007. "DBpedia: A nucleus for a Web of open data," In: Karl Aberer, Key-Sun Choi, Natasha Noy, Dean Allemang, Kyung-Il Lee, Lyndon Nixon, Jennifer Golbeck, Peter Mika, Diana Maynard, Riichiro Mizoguchi, Guus Schreiber and Philippe Cudré-Mauroux (editors). *The semantic Web: Sixth International Semantic Web Conference, Second Asian Semantic Web Conference, ISWC 2007 + ASWC 2007, Busan, Korea, November 11–15, 2007. Proceedings*. Berlin: Springer-Verlag, pp. 722–735.
doi: http://dx.doi.org/10.1007/978-3-540-76298-0_52, accessed 10 November 2016.

Steve Awodey and A.W. Carus, 2009. "From Wittgenstein's prison to the boundless ocean: Carnap's dream of logical syntax," In: Pierre Wagner (editor). *Carnap's logical syntax of language*. Houndmills, England: Palgrave Macmillan, pp. 79–106.

Tim Berners-Lee, 1994. "World Wide Web: Future directions," at <http://www.w3.org/Talks/WWW94Tim/>, accessed 14 July 2016.

Tim Berners-Lee, James Hendler and Ora Lassila, 2001. "The Semantic Web," *Scientific American*, volume 284, number 5, pp. 28–37, and at <https://www.scientificamerican.com/article/the-semantic-web/>, accessed 10 November 2016.

Jacques Bus, Malcolm Crompton, Mireille Hildebrandt and George Metakides, 2012. *Digital enlightenment yearbook 2012*. Amsterdam: IOS Press.

Rudolf Carnap, 2003. *The logical structure of the world; and, Pseudoproblems in philosophy*. Translated by Rolf A. George. Chicago: Open Court.

Rudolf Carnap, 2002. *The logical syntax of language*. Translated by Amethe Smeaton. Chicago: Open Court.

Rudolf Carnap, 1962. *Logical foundations of probability*. Second edition. Chicago: University of Chicago Press; version at http://bcf.usc.edu/~bwrobert/USCPhilSci/docs/Carnap_FoundationsOfProbabilityCh1.pdf, accessed 10 January 2015.

Rudolf Carnap, 1959. "The elimination of metaphysics through logical analysis of language," In: A.J. Ayer (editor). *Logical positivism*. Glencoe, Ill.: Free Press, pp. 60–81.

Rudolf Carnap, 1952. "Meaning postulates," *Philosophical Studies*, volume 3, number 5, pp. 65–73.

A.W. Carus, 2007. *Carnap and twentieth-century thought: Explication as Enlightenment*. Cambridge: Cambridge University Press.

Richard Creath, 2009. "The gentle strength of tolerance: The logical syntax of language and Carnap's philosophical programme," In: Pierre Wagner (editor). *Carnap's logical syntax of language*. Houndmills, England: Palgrave Macmillan, pp. 203–214.

Tony Fry, 2009. *Design futuring: Sustainability, ethics, and new practice*. New York: Berg.

Gottfried Gabriel, 2012. "Carnap, pseudo problems, and ontological questions," In: Pierre Wagner (editor). *Carnap's ideal of explication and naturalism*. Houndmills, England: Palgrave Macmillan, pp. 23–33.

Gottfried Gabriel, 2009. "Carnap's 'Elimination of metaphysics through logical analysis of language,'" *Linguistic and Philosophical Investigations*, volume 8, pp. 53–70.

Nelson Goodman, 1977. *The structure of appearance*. Dordrecht: Reidel.
doi: <http://link.springer.com/10.1007/978-94-010-1184-6>, accessed 20 August 2016.

Nicola Guarino, Daniel Oberle and Steffen Staab, 2009. "What is an ontology?" In: Steffen Staab and Rudi Studer (editors). *Handbook on ontologies*. Berlin: Springer-Verlag, pp. 1–17.
doi: http://link.springer.com/10.1007/978-3-540-92673-3_0, accessed 29 September 2016.

R.V. Guha, Dan Brickley and Steve MacBeth, 2016. "Schema.org: Evolution of structured data on the Web," *Communications of the ACM*, volume 59, number 2, pp. 44–51.
doi: <http://dx.doi.org/10.1145/2844544>, accessed 10 November 2016.

Katie Hafner and Matthew Lyon, 1996. *Where wizards stay up late: The origins of the Internet*. New York: Simon & Schuster.

Donna Haraway, 2016. *Staying with the trouble: Making kin in the Chthulucene*. Durham, N.C.: Duke University Press.

- Patrick Hayes (editor), 2004. "RDF semantics," *W3C* (10 February), at <https://www.w3.org/TR/2004/REC-rdf-mt-20040210/>, accessed 10 November 2016.
- Patrick J. Hayes, 1995. "In defense of logic," In: George F. Luger (editor). *Computation and intelligence: Collected readings*. Menlo Park, Calif.: AAAI Press, pp. 261–273; originally appeared as, Patrick J. Hayes, 1977. "In defense of logic," *IJCAI '77: Proceedings of the Fifth International Joint Conference on Artificial Intelligence*. Volume 1, pp. 559–565.
- Patrick J. Hayes and Peter F. Patel-Schneider (editors), 2014. "RDF 1.1 semantics," *W3C* (25 February), at <https://www.w3.org/TR/rdf11-mt/>, accessed 10 November 2016.
- Edwin Hutchins, 1995. *Cognition in the wild*. Cambridge, Mass.: MIT Press.
- Immanuel Kant, 1963. "What is enlightenment?" In: Immanuel Kant. *On history*. Edited, with an introduction by Lewis White Beck. Translated by Lewis White Beck, Robert E. Anchor and Emil L. Fackenheim. Indianapolis: Bobbs-Merrill, pp. 3–10.
- Leslie Lamport, 1978. "Time, clocks, and the ordering of events in a distributed system," *Communications of the ACM*, volume 21, number 7, pp. 558–565.
doi: <http://dx.doi.org/10.1145/359545.359563>, accessed 10 November 2016.
- Ora Lassila and Ralph Swick (editors), 1999. "Resource Description Framework (RDF) Model and syntax specification," *W3C proposed recommendation* (9 January), at <https://www.w3.org/TR/PR-rdf-syntax/>, accessed 10 November 2016.
- Bruno Latour, 2014. "Agency at the time of the Anthropocene," *New Literary History*, volume 45, number 1, pp. 1–18.
doi: <http://dx.doi.org/10.1353/nlh.2014.0003>, accessed 10 November 2016.
- Michael P. Lynch, 2016. *The Internet of us: Knowing more and understanding less in the age of big data*. New York: Norton.
- John McCarthy, 1962. "Time-sharing computer systems," In: Martin Greenberger (editor). *Management and the computer of the future*. Cambridge, Mass.: MIT Press, pp. 221–236.
- Drew McDermott, 1978. "Tarskian semantics, or no notation without denotation!" *Cognitive Science*, volume 2, number 3, pp. 277–282.
doi: http://dx.doi.org/10.1207/s15516709cog0203_5, accessed 10 November 2016.
- Alexandre Monnin, 2015. "L'ingénierie philosophique de Rudolf Carnap: De l'IA au Web sémantique," *Cahiers philosophiques*, number 141,, pp. 27–53.
doi: <http://dx.doi.org/10.3917/caph.141.0027>, accessed 10 November 2016.
- Thomas Mormann, 2010. "Enlightenment and formal romanticism — Carnap's account of philosophy as explication," In: Juha Manninen and Friedrich Stadler (editors). *The Vienna Circle in the Nordic countries. Vienna Circle Institute Yearbook*, volume 14. Dordrecht: Springer Science + Business Media, pp. 263–279.
doi: http://dx.doi.org/10.1007/978-90-481-3683-4_16, accessed 10 November 2016.
- C. Ulises Moulines, 1991. "Making sense of Carnap's 'Aufbau'," *Erkenntnis*, volume 35, numbers 1–3, pp. 263–286.
doi: http://dx.doi.org/10.1007/978-94-011-3490-3_14, accessed 10 November 2016.
- Douglas Rushkoff, 2010. *Program or be programmed: Ten commands for a digital age*. London: OR Books.
- Paul Arthur Schilpp (editor), 1963. *The philosophy of Rudolf Carnap*. La Salle, Ill.: Open Court.
- Carmela Troncoso, George Danezis, Marios Isaakidis and Harry Halpin, in press. "Systematizing decentralization and privacy: Lessons from 15 years of research and deployments," *Privacy Enhancing Technologies*.
- Thomas Uebel, 2012. "The bipartite conception of metatheory and the dialectical conception of explication," In: Pierre Wagner (editor). *Carnap's ideal of explication and naturalism*. Houndmills, England: Palgrave Macmillan, pp. 117–130.
- Pierre Wagner, 2015. "Carnapian and Tarskian semantics," *Synthese*, pp. 1–23.
doi: <http://dx.doi.org/10.1007/s11229-015-0853-7>, accessed 10 November 2016.
- Pierre Wagner (editor), 2012. *Carnap's ideal of explication and naturalism*. Houndmills, England: Palgrave Macmillan.
- Ludwig Wittgenstein, 1922. *Tractatus logico-philosophicus*. London: K. Paul, Trench, Trubner & Co.
- Terry Winograd and Fernando Flores, 1986. *Understanding computers and cognition: A new foundation for design*. Norwood, N.J.: Ablex Publishing.
- Tim Wu, 2011. *The master switch: The rise and fall of information empires*. New York: Vintage Books.

Received 6 November 2016; accepted 10 November 2016.

Ce(tte) œuvre est mise à disposition selon les termes de la [Licence Creative Commons Attribution — Pas d'Utilisation Commerciale — Pas de Modification 4.0 International](#).

The decentralization of knowledge: How Carnap and Heidegger influenced the Web
by Harry Halpin and Alexandre Monnin.

First Monday, Volume 21, Number 12 - 5 December 2016

<http://firstmonday.org/ojs/index.php/fm/rt/prINTERfriendly/7109/5655>

doi: <http://dx.doi.org/10.5210/fm.v21i12.7109>