

Sense and Reference on the Web

Harry Halpin

Received: date / Accepted: date

Abstract We examine a crucial question for the World Wide Web: What does a Uniform Resource Identifier (URI) mean? Crucial for the next-generation Semantic Web, can it refer to things outside web-pages? The Web as a universal information space for naming and accessing information via URIs. However, the classical philosophical problems of meaning and reference that have been the source of debate within the philosophy of language return when the Web is given as the foundation for a knowledge representation with the Semantic Web. Debates on the Semantic Web about the meaning and referential status of a URI are explored as analogues to debates about the meaning and reference of names in the philosophy of language. Three main positions are inspected: the logical position, as exemplified by the descriptivist theory of reference, the direct reference position, as exemplified by Putnam and Kripke's causal theory of reference, and a Wittgensteinian position that views URIs as a public language, as exemplified by Web search engines. These positions show that debates within the philosophy of language are alive and well on the Web, and so in the philosophy of computer science.

Keywords knowledge representation · philosophy of language · World Wide Web · URI · Semantic Web

Harry Halpin
10 Crichton St.
Edinburgh
EH8 9AB
United Kingdom
Tel.: +44 131 650 4630
Fax: +44 131 651 1426
E-mail: H.Halpin@ed.ac.uk

1 Introduction: The Web and URIs

In the course of the practice of computer science, what appears to be mere disputes over engineering can reveal themselves upon inspection to be fundamentally philosophical problems. Whether or not any new light is cast upon the problems by the impact of computer science is the domain of the philosophy of computer science. However, first we must find these problems. Even the World Wide Web, without a doubt one of the most significant computational phenomena to date, poses fundamentally philosophical problems at its foundation, problems that reflect long-standing debates within the philosophy of language.

What is the World Wide Web? The World Wide Web, often shortened to the 'Web,' consists of a space of names called *Uniform Resource Identifiers* (URIs), a unique identifier whose syntax was invented by Tim Berners-Lee, earning him popular acknowledgement as the 'inventor of the Web'. Examples of URIs include <http://www.google.com> and <http://www.wikipedia.org>. URIs like <http://www.wikipedia.com/wiki/Philosophy> are a separate URI from <http://www.wikipedia.org>, although they both share a *domain name*, the root [wikipedia.org](http://www.wikipedia.org). URIs for usually used for accessing hypertext web-pages. In this regard, the Web can be considered a virtual space for naming information based on URIs built on top of the physical infrastructure of the Internet, the protocols that actually ship the bits that compose the hypertext from the server to the client with the web-browser.

The functioning of the Web can be illustrated by the following example. An agent, such as a web browser, wishes to access some information about a *resource* such as the Eiffel Tower in Paris. A resource is defined in the broadest of terms to be "anything that might be iden-

tified by a URI” (Berners-Lee et al., 2005). Berners-Lee was careful not to define a resource only as information that is accessible via the Web, since not only may resources be “electronic documents” and “images” but also “not all resources are network retrievable; e.g., human beings, corporations, and bound books in a library” (Berners-Lee et al., 1998). A person can access the representation of a resource, such as its web-page, using its URI, like `http://www.tour-eiffel.fr/`. While this is all uncontroversial, the novelty of the Web lays in extending URIs beyond web-pages, as a telephone number can be given a URI such as `tel:+1-816-555-1212`.

These URIs for things outside web-pages remained mostly a theoretical possibility until the advent of the Semantic Web, an ambitious project to the use of the Web itself as the infrastructure for a global knowledge representation system. The essential bet of the Semantic Web, a next generation of the Web “in which information is given well-defined meaning, better enabling computers and people to work in cooperation,” is that decentralized agents will come to an agreement on using the *same* URI to name a thing, including things that aren’t accessible on the Web, like people, places, and abstract concepts (Berners-Lee et al., 2001). Suddenly, the question of what a ‘resource’ is, and how different agents can tell which resources a URI identify, transformed from a purely theoretical question into one with practical consequences for the further development of the Web.

The central question is: What is the meaning of a URI? Our thesis is that URIs on the Web are analogous to a special kind of computational ‘name’ (or to be more precise, a ‘proper name’) and thus the Web and the Semantic Web inherits the classical problems the philosophy of language. We will explore first the Semantic Web project, and outline the centrality of URIs to its success. As began in earlier work Halpin and Thompson (2005) and also delved into in similar terms by Wilks Wilks (2008a), we will outline the debates over the meaning of a URI between Tim Berners-Lee and Pat Hayes as the argument between descriptivist and causal theories of reference in philosophy of language, and outline a third relatively unknown philosophical position inspired by the philosophy of Wittgenstein and the practical success of search engines like Google.

2 The Semantic Web

What makes knowledge representation language on the Semantic Web *different* from classical knowledge representation in computer science? 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 language that has some sort of relatively content-neutral syntax (Berners-Lee, 1994). The parallel to knowledge representation in artificial intelligence is striking, as it also sought to find a syntax for encoding human-intelligence. The second point, of “allowing links,” means that the basic model of the Semantic Web will be a reflection of the Web itself: the Semantic Web is constituted by connecting resources by links (Berners-Lee, 1994).

The Semantic Web is then easily construed as a descendant of semantic networks from symbolic artificial intelligence, where nodes are resources and arcs are links. Semantic networks refer declaratively to things in the world, but uses ‘natural-language-like’ labels on its nodes and edges. Yet semantic networks fell out of favor because of their use of ambiguous natural language terms to identify their nodes and arcs, which became a problem when semantic networks were transported between domains and different users, a problem that would be fatal in the decentralized and multi-lingual environment of the Web (Woods, 1975). When researchers attempted to communicate or combine their knowledge representation schemes, no-one really knew what the natural language description *meant* except the author. As powerfully explained by Woods, the IS-A link in semantic networks was interpreted in at least three different ways (Woods, 1975), which could represent both subclassing, instantiation, close similarity, and more. This led to an assault on semantic networks by champions of first-order logic like Hayes, who believed that by providing a formal semantics that defined ‘meaning’, first-order logic at least allowed knowledge representations to be transportable across domains, and that many alternative knowledge representations could be re-expressed in first order-logic (Hayes, 1977).

Under the aegis of the Web standards body the World Wide Web Consortium (W3C), the *Resource Description Framework* (RDF) was created as the first knowledge representation language for the Semantic Web. It was clearly influenced by work in AI on semantic networks, and this should come as no surprise, for RDF was heavily inspired by the work of Ramanathan V. Guha on the Meta-Content Framework (MCF), which was based on Guha’s earlier experience with semantic networks (Guha, 1996). Before working on MCF, Guha was chief lieutenant of the Cyc project, the most ambitious and philosophically controversial project of classical artificial intelligence (R.V.Guha and D.Lenat,

1993). There are nonetheless some key differences between semantic networks and RDF. First, RDF was supplemented by a formal semantics created by the primary critic of semantic networks, Pat Hayes (Hayes, 2004). Importantly, every component of the knowledge representation language is considered a resource, and thus can be given a URI, replacing what is believed to be ambiguous words in semantic networks with URIs. Since statements in a knowledge representation language are usually about the world outside the Web, this means that the Semantic Web crucially depends on the rather strange fact that URIs can refer to things outside the Web.

The second step in Berners-Lee’s vision for the Semantic Web is “allowing links to be created with relationship values,” (Berners-Lee, 1994). Since RDF is composed of resources, and any resource may link to another resource, then any term in RDF may be linked to another term. This linking forms the heart of RDF, as it allows disparate URIs to be linked together in order for statements in RDF to be made. The precise form of a statement in RDF is a *triple*, which consists of two resources connected by a link, as shown in Figure 1. This use of RDF shows off the flexibility of using URIs and links for reference instead of access.

There are several options for encoding Semantic Web statements. The W3C standardized an encoding of RDF in a verbose XML format called ‘RDF/XML’ and a simpler encoding called *Turtle* for triples (Beckett and Berners-Lee, 2008). In *Turtle*, a triple is three space-delimited terms (the subject, predicate, and object) ended in a period. With `http://www.example.org/` being abbreviated as `ex`, one abbreviates the example in Figure 1 to `ex:EiffelTower ex:hasArchitect ex:Gustave.Eiffel`. As shown in Figure 1, the *only* noticeable difference between RDF and a classical semantic network is the use of URIs to label the nodes and relationships rather than terms from natural language. Similarly, a triple by itself can only state a single assertion, but webs of links may be made between triples. A set of triples that share URIs is called a *graph*, as illustrated in Figure 2 by two triples having the same subject, namely that ‘The Eiffel Tower in Paris has an architect called Gustave Eiffel.’ With the ability to make separate statements using URIs, the main purpose of RDF is revealed to be information integration. Due to their reliance on URIs, RDF graphs can *graph merge*, when two formerly separate graphs combine with each other when they use any of the same URIs. With a common language of URIs, agents can merge information about the resources identified by the URIs in a decentralized manner.

With the ability to make separate statements using URIs, the main purpose of RDF is revealed to be in-

formation integration. Due to their reliance on URIs, RDF graphs can *graph merge*, when two formerly separate graphs can combine with each other when they use any of the same URIs.

So the only substantive difference between traditional knowledge representation and the Semantic Web is the central role of URIs. The true bet of the Semantic Web is *not* on the return of symbolic knowledge representation languages. As this use of URIs as the basic element of meaning is central to the Semantic Web, and as it is a genuinely *new* technical claim, it is precisely in the understanding of the status of meaning and reference of URIs that any new *philosophical* claim could possibly be made.

3 The Identity Crisis

How can agents determine what a URI identifies? To use a word more familiar to philosophers, how can anyone determine what a URI refers to or means? On the pre-Semantic Web, URIs trivially identify the hypertext web-pages that those URI allow access to. On the Semantic Web, a whole new cluster of questions, dubbed the *Identity Crisis*, emerges. Can a URI for the Eiffel Tower be used to refer to the Eiffel Tower in Paris itself? If one just re-uses a URI for a web-page of the Eiffel Tower, then one risks the URI being ambiguous between the Eiffel Tower itself and a particular representation of the Eiffel Tower. If one gives the Eiffel Tower *qua* Eiffel Tower its own URI, should that URI allow access to any information, such as a hypertext web-page? This cluster of questions has been dubbed the *Identity Crisis* of the Semantic Web.

As regards any theory of meaning for URIs, in the realm of official Web standards, the jury is still out. In the specification of RDF, Hayes notes that “exactly what is considered to be the ‘meaning’ of an assertion in RDF or RDF(S) in some broad sense may depend on many factors, including social conventions, comments in natural language” so unfortunately “much of this meaning will be inaccessible to machine processing” such that a “a full analysis of meaning” is “a large research topic” (Hayes, 2004). Unsurprisingly, the reason there is no standardized way to determine the meaning of a URI is because, instead of a single clear answer, there is a conceptual quagmire dominated by two positions.

The first position, the *direct reference position*, is that the meaning of a URI is whatever was intended by the owner. The owner of the URI should be able to unambiguously declare and communicate the meaning of any URI, including a Semantic Web URI. In this position, the referent is generally considered to be some individual unambiguous *single* thing, like the Eiffel Tower

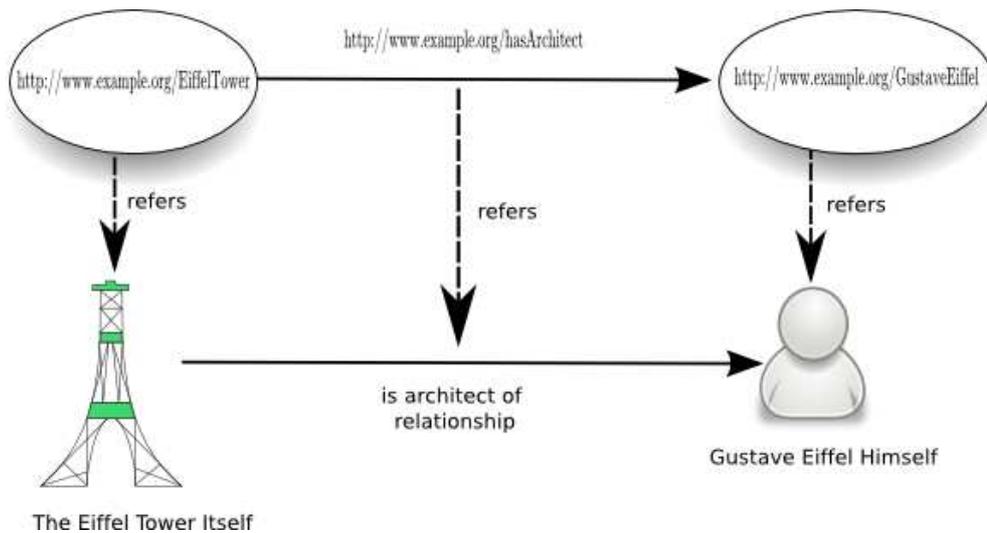


Fig. 1 An example RDF statement

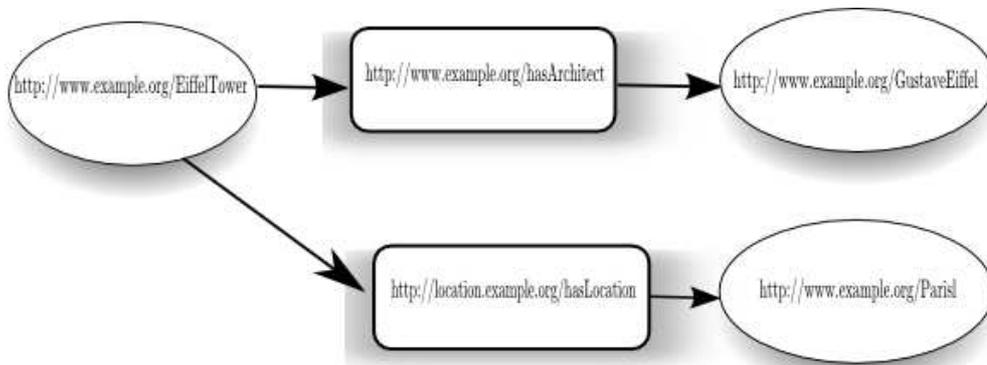


Fig. 2 Merging RDF triples

or the concept of a unicorn. This viewpoint is the one generally held by many Web architects, like Berners-Lee, who imagine it holds not just for the Semantic Web, but the entire Web.

The second position, which we call the *logical position* due to its more clear roots in non-modal logic, is that for the Semantic Web, the meaning of a URI is given by whatever things satisfy the model(s) given by the formal semantics of the Semantic Web. Adherents of this position hold that the referent of a URI is ambiguous, as many different things can satisfy whatever model is given by the interpretation of some sets of sentences using the URI. This position is generally held by logicians, who claim that the Semantic Web is entirely distinct from the hypertext Web.

These two antagonistic positions were subterranean in the development of the Semantic Web, until a critical point was reached in an argument between Pat Hayes, the AI researcher primarily responsible for the formal

semantics of the Semantic Web, and Berners-Lee. During the creation of RDF, Berners-Lee then put forward his own viewpoint that “a single meaning is given to each URI,” which is summarized by the slogan that a URI “identifies one thing.” (Berners-Lee, 2003b). In response, Hayes said that “it is simply untenable to claim that all names identify one thing” (Hayes, 2003a). Furthermore, he goes on to state that this is one of the basic results of the knowledge representation community and 20th century linguistic semantics, and so that the W3C cannot by fiat render the judgment that a URI identifies one thing. Berners-Lee rejects Hayes’s claim that the Semantic Web must somehow build upon the results of logic and natural language, instead claiming that “this system is different from natural language: we designed it such that each URI identifies one and only one concrete thing in the real world or one and only one globally shared concept” (Berners-Lee, 2003a). In exasperation, Hayes retorted that “I’m not saying that the ‘unique

identification’ condition is an unattainable ideal: I’m saying that it doesn’t make sense, that it isn’t true, and that it could not possibly be true. I’m saying that it is *crazy*” (Hayes, 2003b). One should be worried when two researchers such as Berners-Lee and Hayes have such a titanic disagreement, where no sort of consensus agreement seems forthcoming. What is remarkable is that the position of Hayes corresponds to a well-known theory of meaning and reference, the ‘descriptivist theory of reference’ attributed to early Wittgenstein, Carnap, Russell, and given its most precise logical form by Tarski (Luntley, 1999). However, it is common currency in philosophical circles that the descriptivist theory of reference was seriously challenged by the ‘causal theory of reference’ championed by Kripke and extended by Putnam (Luntley, 1999). It is precisely this causal theory of reference that Berners-Lee justifies in his direct reference position. Thus, the curious coincidence is that both opposing positions on the Semantic Web correspond to *equally* opposing positions in philosophy. This correspondence was also noticed by Wilks, but we give much more detail and come to differing conclusions Wilks (2008a).

4 The Logical Position and the Descriptivist Theory of Reference

The origin of the logical position is in what is popularly known as ‘the descriptivist theory of reference’. In the *descriptivist theory of reference*, the referent of a name is given by whatever satisfies the descriptions associated with the name. Usually, the descriptions are thought to be logical statements, so a name is actually a disguised logical description. The referent of the name is then equivalent to the set of possible things, given normally by a mathematical model, such that all statements containing the name are satisfied.

4.1 Logical Atomism

The roots of the descriptivist theory of reference lay with the confluence of philosophers who are known as *logical atomists*, a term coined by Bertrand Russell, and influential to later epistemological projects like the logical positivism of Rudolf Carnap. Although eventually abandoned by Bertrand Russell, logical atomism is a vast school of thought that has proven tremendously influential, even in its current discredited state, for our purposes we will only concern ourselves with one particular doctrine: The problem of how natural language terms relate to the logical descriptions, and logical descriptions to the world.

Bertrand Russell begins the investigation of the connection between logic and language is his landmark investigation *On Denoting* with a deceptively simple question: “is the King of France bald?” (Russell, 1905). To what referent does the description “the King of France” refer to? (Russell, 1905) Since in Russell’s time there was no King of France, it could not refer to anything like what Carnap later called “elementary sense data” (Carnap, 1928). In this regard, Russell makes a crucial distinction. According to Russell, elementary sensory experiences are known through *acquaintance*, in which we have some sort of direct ‘presentation of’ the thing (Russell, 1905). According to Russell, these statements of acquaintance with directly present sensory data employ what are known as Russellian demonstratives (such as ‘this’ or ‘that’) as exemplified by the statement “That is the Eiffel Tower.” Yet knowledge of a thing can be based on *description*, which are those “things we only reach by means of denoting phrases” (Russell, 1905). Russell believed that “all thinking has to start from acquaintance, but it succeeds in thinking *about* many things with which we have no acquaintance” via the use of description (Russell, 1905). Russell was most interested in whether those things with which we have direct acquaintance can be considered true or false, or whether a more mysterious third category such as ‘nonsense’ is needed. Russell opts to reject creating imaginary but true ‘things’ as well as any third category, but instead holds that statements such as “the King of France is bald” are false, since “it is false that there is an entity which is now the King of France and is bald” (Russell, 1905). This solution then raises the alarming possibility that “the King of France is not bald” may also come out false, which would seem to violate the Law of the Excluded Middle. So, Russell counters this move by introducing the fact that “the King of France is bald” is actually a complex logical statement involving scope and quantification, namely $(\exists x.F(x) \wedge G(x)) \wedge (\forall y.F(y) \rightarrow x = y)$, where F is “being the King of France” and G is “being bald” (Russell, 1905). According to the analysis, ‘The King of France’ is merely a *disguised* complex logical statement. Furthermore, this treatment can be extended to proper names such as ‘Sir Walter Scott,’ who can be identified with ‘the author of Waverly,’ so that instead of being a tautology, even a proper name of a person, even if known through acquaintance, is sort of short-hand for a large cluster of logical statements. To use our previous example, the ‘Eiffel Tower’ can be thought of as a short-hand for not only that ‘there exists an entity known as the Eiffel Tower’ but also the logical statement was ‘the aforementioned entity had Gustave Eiffel as its architect.’ If someone did not know that ‘the aforementioned

entity was also the tallest building in the world up until 1930,’ one could then make a statement such as ‘The Eiffel Tower is identical to the tallest building in the world up until 1930’ without merely stating a tautology, and such a statement would add true and consistent knowledge to a hearer who was not aware of the statement.

4.2 Tarski’s Formal Semantics

Tarski abandoned the quaint epistemology of logical atomism in terms of direct acquaintance with sensory data and defined reference purely in terms of mathematical logic in his *The Concept of Truth in Formal Languages* (Tarski, 1935). Reference was just defined as a consequence of the truth *only* in terms of satisfaction of a formal language (Tarski, 1935). To set up his exposition, Tarski defines two languages, the first being the syntactic *object language* L and the second being the meta-language M . The *meta-language* should be more expressive such that it can describe every sentence in the object language, and furthermore, that it contain axioms that allow the truth of every sentence in the object language to be defined. In his first move, Tarski defines the formal conception of truth as ‘Convention T,’ namely that for a given sentence s in L , there is a statement p in M that is a theorem defining the truth of s , that is, the truth of s is determined via a translation of s into M (Tarski, 1935). Tarski then later shows that truth can be formally defined as “ s is true if and only if p ” (Tarski, 1944). For example, if the object language is exemplified by a sentence uttered by some speaker of English and the meta-language was an English description of the real world; ‘The Eiffel Tower is in Paris’ is true if and only if the Eiffel Tower is in Paris. The sentence ‘The Eiffel Tower is in Paris’ must be satisfied by the Eiffel Tower *actually being* in Paris. While this would at first seem circular, its non-circularity is better seen through when the object language is not English, but another language such as German. In this case, “‘Der Eiffelturm ist in Paris’ is true if and only if the Eiffel Tower is in Paris.” However, Tarski was not interested in informal languages such as English, but in determining the meaning of a new formal language via translations to mathematical models or other formal languages with well-known models. If one was defining a formal semantics for some fragment of a knowledge representation language like RDF, a statement such as `http://www.eiffeltower.example.org ex:location ex:Paris` is true if and only if $\exists ab.R(a,b)$ where R , a , and b are given in first-order predicate logic.

If one is defining a formal Tarski-style semantics for a language, what should one do when one encounters

complex statements, such as ‘the Eiffel Tower is in Paris and had as an architect Gustave Eiffel’? The answer is at the heart of Tarski’s project, the second component of Tarski’s formal semantics is to use the principle of compositionality so that any complex sentence can have its truth conditions derived from the truth conditions of its constituents. To do this, the meta-language has to have finitely many axioms, and each of the truth-defining theorems produced by the meta-language have to be generated from the axioms (Tarski, 1935). So, the aforementioned complex sentence is true if and only if $\exists ab.R(a,b) \wedge Q(a,c)$, where Q can be the *architect of* relationship, c can be Gustave Eiffel and a the Eiffel Tower. Tarski’s theory as explained so far only deals with ‘closed’ sentences, i.e. sentences containing no variables or quantification. The third, and final component of Tarski’s formal semantics is to use the notion of satisfaction via extension to define truth (Tarski, 1935). For a sentence such as ‘all monuments have a location,’ we can translate the sentence to $\forall a.l.monument(a) \rightarrow hasLocation(a,l)$ which is true if and only if there is an extension x from the world that satisfies the logical statements made about a . In particular, Tarski has as his preferred extensions infinite ordered pairs, where the ordered set could be anything (Tarski, 1935). For formal languages, a model-theoretic semantics with a model composed by set theory was standard. For example, the ordered pairs in some model of *(EiffelTower, Paris)* would satisfy, as would *(ScottMonument, Edinburgh)* but not *(Paris, EiffelTower)*. However, there is no reason why these models could not be “God Forthcoming,” things in the the real world itself, albeit given in set-theoretic terms (Smith, 1995). To summarize Tarski’s remarkably successful programme, model-theoretic semantics can produce a theory of truth that defines the semantics of a sentence in terms of the use of a translation of the sentence into some formal language with a finite number of axioms, then using compositionality to define the truth of complex sentences in terms of basic sentences, and finally determining the truth of those basic sentences in terms of what things in a model satisfy the extensions of the basic sentences as given by the axioms. This work marks the high-point of the logical programme, as all questions of meaning are reduced to questions about giving the interpretation of a sentence in terms of a formal notion of truth. This notion of truth is not restricted by the logical atomist’s epistemology of elementary sense data, but instead can range over any possible formal language and any possible world. This victory is not without its costs, since while Tarski provides the best account of the relationship between logical descriptions and the world by simply removing all questions that cannot be phrased formally, formal se-

mantics by itself leaves unsolved the fundamental question about how natural language relates to our sensory experience of the world. Ignoring a problem does not make it go away. So when confronted with this vexing problem, champions of formal semantics often revert to the Russellian doctrine of direct acquaintance, as demonstrated in the next section.

4.3 Logical Descriptions Unbound on the Web

While the descriptivist theory of reference seems distant from the Identity Crisis of the Web, it is in fact central to the position of Hayes and the Semantic Web as a whole. This is primarily because Hayes’s background was in the logical tradition, with his particular specialty being the creation of Tarski-style semantics for knowledge representation languages. What Hayes calls the “basic results in 20th century linguistic semantics” that Berners-Lee’s dictum that “URIs identify one thing” violates is the interpretation of URIs in a Tarski-style formal semantics (Hayes, 2003a). For the logical position, the *semantics* in the Semantic Web derive from the Tarski-style formal semantics Hayes created for the Semantic Web (Hayes, 2004).

Before delving into the formal semantics of RDF, it should be noticed that these semantics are done by extension, like most other formal languages Hayes (2004). However, the semantics of RDF are purposefully quite weak as not to allow arithmetic or constructs like the negation of a class, and so RDF avoids logical paradoxes like the encoding of Gödel sentences. Yet in order to make RDF triples as flexible as possible, RDF includes features normally associated with higher-order logic such as “a property may be applied to itself” and classes “may contain themselves” (Hayes, 2004). This is handled semantically by having first an interpretation map the URI to an individual. Then unlike standard first-order logic, this individual then maps to different extensions depending on the role the URI is playing as a property or class in the triple. A simple example should suffice to give a flavour of the formal semantics, where a relation is just another kind of individual. What is the formal semantics of `ex:EiffelTower ex:architect ex:Gustave_Eiffel`? To simplify slightly, Hayes defines the formal semantics in terms of set theory, where there is a set of resources that compose the model of the language, a set of properties, and a set of URIs that can refer to resources. The interpretation of any RDF statement is then given as an extensional mapping from the set of properties to the powerset of resources, to the set of pairs of resources. So, given a set-theoretic model consisting of elements (given by italics) *Gustave Eiffel* and *the Eiffel Tower* and *being the architect of*, then

`ex:EiffelTower` \models *the Eiffel Tower*, `ex:Gustave_Eiffel` \models *Gustave Eiffel* and `ex:architect` \models *being the architect of*, so that the entire triple maps to a set of pairs: `ex:EiffelTower ex:architect ex:Gustave_Eiffel` \models $(\dots, (\textit{the Eiffel Tower}, \textit{Gustave Eiffel}), \dots)$. Common-sense human intuitions will likely have this interpretation maps to `ex:EiffelTower ex:architect ex:Gustave_EiffelTower`, and using the axioms defined in the RDF formal semantics, a few new triples can be inferred, such as `ex:architect rdf:type rdf:Property`, i.e. *being an architect of* is a property of something.

However, the inherent pluralism of the Tarski approach to models also means that another equally valid interpretation would be the inverse, i.e. the mapping of `ex:EiffelTower` to *Gustave Eiffel* and `ex:Gustave_Eiffel` to *the Eiffel Tower*. In other words, `ex:architect` \models *being the architect of*, so that the entire triple maps to a set of pairs `ex:EiffelTower ex:architect ex:Gustave_Eiffel` \models $(\dots, (\textit{Gustave Eiffel}, \textit{Eiffel Tower}), \dots)$. Due to the unconstrained nature of RDF, `ex:architect` has no ‘natural’ relationship to anything in particular, but could easily be assigned either *the Eiffel Tower* or *Gustave Eiffel* just as easily as *being the architect of*. Furthermore, the model could just as easily be given by something as abstract as the integers 1 and 2, and an equally valid mapping would be for `ex:EiffelTower` \models 1 and `ex:Gustave_Eiffel` \models 2, so that `ex:architect` \models *being the architect of*, so that the entire triple maps to a set of pairs `ex:EiffelTower ex:architect ex:Gustave_Eiffel` \models $(\dots, (1, 2), \dots)$. Indeed, the extreme pluralism of a Tarski-style semantics shows that, at least if all one has is a single lone triple statement, that triple can be satisfied by any model. This is no mere oddity of formal languages, this would also hold for any lone sentence in a language like English – such as “Gustave Eiffel is the architect of the Eiffel Tower” – as long as one subscribed to a Tarski-style semantics for natural language. As the number of triples increased, the amount of possible things that satisfy the model is thought to decrease, but in such a loose language as RDF, Hayes notes that it is “usually impossible to assert enough in any language to completely constrain the interpretations to a single possible world, so there is no such thing as ‘the’ unique interpretation” (Hayes, 2004). This descriptivist theory of reference, where descriptions are logical statements in RDF, is illustrated in Figure 3.

While Hayes makes no claim that access to some web-pages via URIs is not possible, he claims that such access to Web representations is orthogonal to the question of what a URI could refer to, since “the architecture of the Web determines access, but has no direct influence on reference” (Hayes and Halpin, 2008). Furthermore, Hayes’s logical understanding of ambigu-

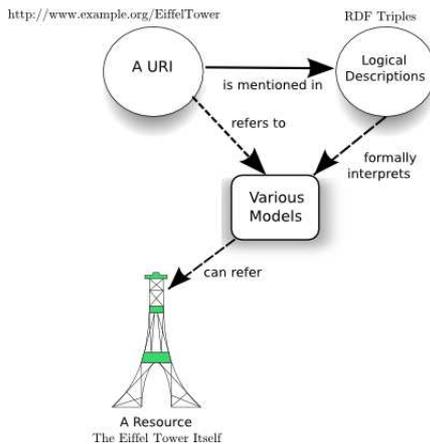


Fig. 3 The descriptivist theory of reference for URIs

ity parts path with natural language understandings of ambiguity: Hayes claims that reference to resources is completely *independent* of whatever Web representations can be accessed, even if those contain logical expressions. While much credit should be given to Hayes for creating a logical semantics for RDF, the problem of connecting these descriptions to the world outside of the Web falls outside formal semantics and so opens up a seemingly uncrossable abyss between the logical descriptions and sensory data. One seemingly easy way out of this abyss is to revert to the doctrine of Russellian direct acquaintance, also known as ostension. In moments, Hayes himself seems to subscribe to the logical atomist epistemology of Russell, as he says that “reference can either be established by either description or ostention” with ostention being defined as the use of a Russellian demonstrative (like ‘that’ or ‘this’) identifying a particular “patch of sense data” via a statement such as ‘that is the Eiffel Tower’ (Hayes, 2006). Since most of the things referred to by names are not accessible, reference can only be determined by description, and these descriptions are inherently ambiguous as regards any sense data (Hayes and Halpin, 2008).

As our example showed, RDF in general says so little inferentially that many different models can satisfy almost any given RDF statement. Therefore, Hayes considers it essential to ditch the vague word ‘identify’ as used in URIs, and distinguish between the ability of URIs to access and refer. While access is constrained by Web architecture, according to Hayes, reference is absolutely unconstrained except by formal semantics, and so “the relationship between access and reference is essentially arbitrary” (Hayes and Halpin, 2008). From this philosophical position, the Identity Crisis dissolves into a pseudo-problem, for the same URI can indeed access a web-page and refer to a person unproblematically, as

they no longer have to obey the dictum to identify one thing. Hayes compares this situation to that of *overloading*, using a single name to refer to multiple referents, and instead of being a problem, “it is a way of using names efficiently” and not a problem for communication, as “natural language is rife with lexical ambiguity which does not hinder normal communication,” as these ambiguities can almost always be resolved by sufficient context (Hayes and Halpin, 2008). Overall, the argument of Hayes against Berners-Lee in the Identity Crisis is the position of keeping the formal semantics of reference separate from the Web.

5 The Direct Reference Position and the Causal Theory of Reference

The alternative slogan of Berners-Lee, that “URIs identify one thing,” may not be completely untenable after all (Berners-Lee, 2003b). It appears to even be intuitive, for when one says ‘I went to visit the Eiffel Tower,’ one believes one is talking about a very *particular* thing in the *real* world called the ‘Eiffel Tower,’ not a cluster of descriptions or model of the world. The direct theory of reference of Berners-Lee has a parallel in philosophy, namely Saul Kripke’s ‘causal theory of reference,’ the widely-known argument against the descriptivist theory of reference, and so the position of Hayes (Kripke, 1972). In contrast to the descriptivist theory of reference, where the content of any name is determined by ambiguous interpretation of logical descriptions, in the *causal theory of reference* any name refers via some causal chain directly to a referent (Kripke, 1972).

5.1 Kripke’s Causal Theory of Proper Names

The causal theory of reference was meant to be an attack on the descriptivist theory of reference attributed to Russell, and its effect in philosophy has been to discredit any neo-Russellian descriptivist theory of reference (Luntley, 1999). Unsurprisingly, the causal theory of reference also has its origin in logic, since Kripke as a modal logician felt a theory of reference was needed that could make logical statements about things in different logically possible worlds (Kripke, 1972). However, while Kripke did not directly confront the related position of Tarski, his argument does nonetheless attempt to undermine the ambiguity inherent in Tarski’s model-theoretic semantics, although a Tarski-style semantics can merely ‘flatten’ models of possible worlds into a singular model (Luntley, 1999). Still, as a response in philosophy of language, it is accepted as a classical refutation of the descriptivist theory of reference.

In Kripke's *Naming and Necessity*, an agent fixes a name to a referent by a process called *baptism*, in which the referent, known through direct acquaintance is associated with a name via some local and causally effective action by the agent (Kripke, 1972). Afterwards, a historical and causal chain between a current user of the name and past users allows the referent of a name to be transmitted unambiguously through time, even in *other possible worlds*. For example, a certain historical personage was given the name 'Gustave Eiffel' via a rather literal baptism, and the name 'Gustave Eiffel' would still refer to that baptized person, even if he had not been the architect of the Eiffel Tower, and so failed to satisfy that definite description. Later, the causal chain of people talking about 'Gustave Eiffel' would identify that very person, even after Gustave Eiffel was dead and gone. Descriptions aren't entirely out of the picture on Kripke's account; they are necessary for disambiguation when the context of use allows more than one interpretation of a name, and they figure in the process by which things actually get their names, if the thing cannot be directly identified. However, this use of descriptions is a mere afterthought with no causal bearing on determining the referent of the name itself, for as Kripke puts it, "let us suppose that we do fix the reference of a name by a description. Even if we do so, we do not then make the name synonymous with the description, but instead we use the name rigidly to refer to the object so named, even in talking about counterfactual situations where the thing named would not satisfy the description in question" (Kripke, 1972). So what is crucial is not satisfying any description, but the act of baptism and the causal transmission of the name.

5.2 Putnam's Theory of Natural Kinds

Kripke's examples of the causal theory of reference used proper names, such as 'Cicero' or 'Aristotle,' and he did not extend his analysis to the whole of language in a principled manner. However, Hilary Putnam, in his *The Meaning of 'Meaning,'* extends Kripke's analysis to all sorts of names outside traditional proper names, and in particular Putnam uses for his examples the names of natural kinds (Putnam, 1975). Putnam was motivated by an attempt to defeat what he believes is the false distinction between intension and extension. The set of logical descriptions, which Putnam identifies with a "psychological state," that something must satisfy to be given a name is the *intension*, while those things in a given interpretation that actually satisfy these descriptions, is the *extension* (Putnam, 1975). Putnam notices that while a single extension can have multiple

intensions it satisfies, such as the Eiffel Tower both being "in Paris" and "a monument," a single intension is supposed to have the same extension in a given interpretation. If two people are looking for a "monument in Paris," the Eiffel Tower should satisfy them both, even though the Eiffel Tower can also have many other possible descriptions.

Putnam's analysis can be summarized as follows: Imagine that there is a world "very much like Earth" called 'Twin Earth.' On Twin Earth "the liquid called 'water' is not H_2O but a different liquid" whose chemical formula is abbreviated as XYZ , and that this XYZ is "indistinguishable from water at normal temperatures and pressures" since it "tastes like water and quenches thirst like water" (Putnam, 1975). A person from Earth would *incorrectly* identify XYZ for their normal referent of water, as it would satisfy all their descriptions. In this regard, this shows that meanings "ain't in the head" but are in fact determined, not by individual language use or descriptions, but by some indexical relationship to "stuff that is like water around here" normally. That "stuff" *should* get its name and meaning from *experts*, since "probably every adult speaker even knows the necessary and sufficient condition 'water is H_2O ,' but only a few adult speakers could distinguish water from liquids which superficially resembled water...in case of doubt, other speakers would rely on the judgment of these 'expert' speakers" who would ideally test XYZ and determine that it was indeed, not water" (Putnam, 1975). Indeed, less outlandish examples, such as the difference between "beech trees" and "elm trees" are trotted out by Putnam to show that a large amount of our names for things, perhaps even extending beyond natural kinds, are actually determined by expert knowledge (Putnam, 1975). In this way, Kripke's baptism can extend to almost all languages, and scientists can be considered a special sort of naming authority capable of baptizing all sorts of things with a greater authority than everyone else. As even Putnam explicitly acknowledges "Kripke's doctrine that natural-kind words are rigid designators and our doctrine that they are indexical are but two ways of making the same point" (Putnam, 1975).

5.3 Direct Reference on the Web

This causal theory of reference is naturally close to the direct reference position of Berners-Lee, whose background is in expert-created databases. He naturally assumes the causal theory of reference is uncontroversial, for in database schemas, what a term *refers to* is a matter best left to the expert designer of the database. So Kripke and Putnam's account of unambiguous names

can then be transposed to the Web with a few minor variations in order to obey Berners-Lee’s “crazy” dictum that “URIs identify one thing” regardless of interpretation or even accessible web-page (Berners-Lee, 2003b). While it may be a surprise to find Berners-Lee to be a closet Kripkean, Berners-Lee says as much, “that the Web is not the final arbiter of meaning, because URI ownership is primary, and the look-up system of HTTP is...secondary” (Berners-Lee, 2003b). There is also an element of Grice in the direct theory of reference, for the *intended* interpretation and perhaps even purpose of the owner is the one that really matters to Berners-Lee, not any publicly accessible particular Web representation (Grice, 1957). However, ultimately Berners-Lee has far more in common with the causal theory of reference, since although the URI owner’s intention determines the referent, after the minting of the new URI for the resource, the intended interpretation is somehow never supposed to vary (Berners-Lee, 1998).

To apply the causal theory of reference as to URIs, baptism is given by the registration of the domain names, which gives a domain name and legally binding set of IP addresses, such as *example.org*, a legally binding owner. Of course, the natural question then would be if this Kripkean practice can then be extended to entire URIs such as *http://www.example.org/Eiffel*? For most domain names a specific policy given by the owner could set the allowed referents for the creation of URIs that involve the domain name in question, perhaps as embodied in some software system. One could imagine several variations on this theme, from the URIs being controlled indirectly by systems-programmers or even outsourced to the general public in the form of a user-generated URI registry with a single top-level domain. Regardless of the details, the referent of a URI is established by fiat by the owner(s), and then optionally can be communicated to others in a causal chain in the form of publishing web-page accessible from the URI or by creating Semantic Web statements about the URI. This causal theory of reference for URIs is illustrated in Figure 4.

In this manner, the owner of the URI can thereby determine the referent of the URI and communicate it to others, but ultimately the act of baptism and so the determination of the referent are in the hands of the owner of the URI, the self-professed ‘expert’ in the new vocabulary term introduced to the Semantic Web by his URI, and the owner has no real responsibility to host any Web representations at the URI. Since the owner can causally establish a name for a non-Web accessible thing via simply minting a new URI without hosting *any* web-page, under the causal theory of reference the Semantic Web can be treated as having a giant transla-

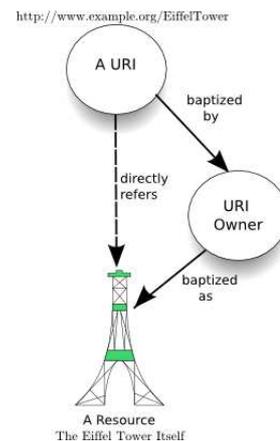


Fig. 4 The causal theory of reference for URIs

tion manual mapping URIs directly to referents, where the URIs refer directly to objects in the world outside of the Web. Realistically, if an agent got a URI like *http://www.example.org/Gustave_Eiffel* and one wanted to know what the URI referred to, one could use a service such as *whois* to look up the owner of the URI, and then contact the owner of the URI if there was any doubt in the matter. Yet since obviously such URIs cannot access things outside the Web and contacting the owner every time a URI is to be used is absurd, what kinds of web-pages, if any, should this giant Semantic Web dictionary return? If it returns no web-page, how can a user-agent distinguish a URI for a referent outside the Web from that of a URI for a web-page? This question is partially answered by Berners-Lee in a solution called ‘303 redirection,’ where a distinct URI is given to the thing-in-of-itself, and then when this URI is accessed by an agent such as a web-browser, a particular Web mechanism called the 303 Header redirects to the agent to another URI for a web-page describing the resource, either in RDF or in HTML, or both. However, this mechanism has been considered difficult to use and understand, “analogous to requiring the postman dance a jig when delivering an official letter” (Hayes, 1977).

6 Sense and Reference Redux

The Semantic Web has still not experienced the tremendous growth of the hypertext Web, and the primary reason appears to be this impasse at the Identity Crisis. For the first few years of its existence (2001-2006), in general the arguments of Hayes prevailed, and the URIs used in RDF graphs did not access any web-pages. However, in this phase of its existence, the Semantic Web did not progress beyond yet another little-used knowledge representation language. In the last few

years (2006-2009), the Semantic Web has experienced phenomenal growth under the term ‘Linked Data,’ as Berners-Lee’s position has had more acceptance and users have started deploying RDF using actual URIs. This growth of estimated billions triples, including large-scale projects by biomedical community and in government data in using the Semantic Web, seems to have implicitly validated Berners-Lee’s direct reference position. Yet that is far from true; what is apparent from any analysis of the Semantic Web is that there appear to be *too many* URIs for some things, while *no* URIs for other things (Halpin and Lavrenko, 2009). As differing users export their data to the Web in a decentralized manner, new URIs are always minted, and so running the risk of fracturing the Semantic Web into isolated ‘semantic’ islands instead of becoming a unified web, as the same URIs are not re-used. The critical missing element of the Semantic Web is some mechanism that allows users to come to agreement on URIs and then share and re-use them, a problem ignored both by the logical and direct reference positions.

The philosophical root of the problem may be that both Russell and Kripke - and so both Berners-Lee and Hayes - reject the notion of ‘sense.’ The Fregean distinction between ‘sense’ and ‘reference’ that provoked both Russell and Kripke’s intellectual projects to build an entire theory of meaning on top of only reference, where Frege held that the the meaning of any term in a language is determined by what Frege calls the “sense” of the sentences that use the term, rather than any direct reference of the term (Frege, 1892). Frege defines ‘sense’ in terms of the mysterious *mode of presentation*, for “to think of then being connected with a sign (name, combination of words, letters), besides that to which the sign refers, which may be called the reference of the sign, also what I should like to call the sense of the sign, wherein the mode of presentation is contained” (Frege, 1892). This rather cryptic statement has caused multiple decades of debate by philosophers of language. The reason the notion of sense was thought of as so objectionable by many philosophers like Russell and Kripke was that sense was viewed as a mysterious notion, perhaps a private and individual notion, much like the Lockean notion of an *idea*. However, Frege himself clearly rejects this, strictly separating the notion of a sense from an individual subjective idea of a referent. Far from subjective, Frege believed that sense was inherently *objective*, “the reference of a proper name is the object itself which we designate by using it; the idea which we have in that case is wholly subjective, in between lies the sense, which is indeed no longer subjective like the idea, but is yet not the object itself” (Frege, 1892). A sense is objective insofar as it is

a shared part of an inherently public language, since a sense is the “common property of many people, and so is not a part of a mode of the individual mind. For one can hardly deny that mankind has a common store of thoughts which is transmitted from one generation to another” (Frege, 1892). While the exact nature of a sense is still unclear, its main characteristic is that it should be whatever is *objectively shared* between agents as regards their use of terms in a language.

It is precisely this notion that sense is ‘objective’ that allows us to construct a new position. Yet how does this notion of sense play out in the philosophy of computer science? Dummett provides an insightful hint, “Frege’s thesis that sense is objective is thus implicitly an anticipation of Wittgenstein’s doctrine that meaning is use” (Dummett, 1993). So we must outline a third position, the public language position that takes the objective notion of ‘sense’ and Wittgenstein’s analysis of “meaning as use” as its foundation, and one that is rooted in the actual use of URIs in computer science, namely in Web-scale information retrieval (Wittgenstein, 1953) .

7 The Public Language Position

It is precisely the social and so public notion of language that has been strangely missing from the debates on reference and meaning on the Semantic Web so far. Wittgenstein points to a social and public notion of language when he says, “Do not ask yourself ‘how does it work with me?’ – ask ‘What do I know about someone else?’” (Wittgenstein, 1953). A language is public and inexorably social, involving more than one agent. So a third position, in contrast to both the logical and direct reference positions, can now be staked. The *public language position* states that since the Web, including the Semantic Web, is a form of language, as a language exists as a mechanism for co-ordination among multiple agents, then the meaning of a URI is the use of the URI by a community of agents.

To contrast this position with the direct reference position, the meaning of a URI is not determined by whatever referent is assigned to it by its owner, unless the owner and other agents actually can come to an agreement on its meaning. The public language position does not give the owner of a URI any particular privilege, except for the obvious asymmetric technical privilege of having the ability to influence the use of the URI through hosting an accessible web-page or redirecting to another URI.

Just as we found analogues of the descriptivist and causal theories of reference in the position of Hayes and Berners-Lee, is there is existing computational analogue

of this third public language position? We believe so, and the answer lies in the theory of meaning implicit in information retrieval: the derivat of meaning via massive statistics that approximate the ‘use’ of a URI on the Web.

Unlike the causal theory of reference and the descriptivist theory of reference, Wittgenstein does not equate the meaning of a sentence with truth or the satisfaction of a model unless these are accepted by all agents that use the URI. Wittgenstein retorts that only “in our language” can “we apply the calculus of truth” (Wittgenstein, 1953). However, we would hold that Wittgenstein is being too curt, and that logical-based inference can approximate the shared meaning between agents to a limited extent.

7.1 Wittgenstein against Private Language

One of the hidden presumptions of the descriptivist and causal theory of reference is the tradition that language can be a *private* phenomenon, that it can be possessed and used by a single agent to accurately describe and refer to the world. Wittgenstein, whose *Tractatus* was the original inspiration for this position, returned to refute this point in his *Philosophical Investigations* (Wittgenstein, 1953). Wittgenstein attacks the very idea of a private language, a language that is somehow only understood by a single person and hence untranslatable to other languages, where “the individual words of this language are to refer to what can only be known to the person speaking; to his immediate private sensations. So another person cannot understand the language” (Wittgenstein, 1953). His primary example is the use of a language to describe sensations of pain. Wittgenstein argues that such a language is absurd, as there would be no “right” way to use the private word for the sensation, for “whatever is going to seem right to me is right” (Wittgenstein, 1953). In his second famous attack on private language, Wittgenstein phrases an attack on private codes of behavior in the infamous example of rule-following in a game like chess, stating that, “It is not possible that there should have been only one occasion on which only one person obeyed a rule” (Wittgenstein, 1953). There can be no norms for behavior, and therefore no meaning, in a private language. This follows from the insight that norms ultimately involve *others*, where the norm is repeated in different circumstances and mediates the collective behavior of multiple agents.

Strangely enough, there is a deep affinity between both the descriptivist and causal theories of reference, for a Kripkean baptism is just some sort of *ostentative* relationship between sense data and a name, ex-

emplified by the act of saying ‘the name of that is the Eiffel Tower.’ This account of baptism is precisely the same as Russell’s account of the use of names via direct acquaintance. A Russellian descriptivist would simply have some ‘sense-data’ that they could label with ‘that is an iron tower’ and then generalize to other sets of ‘sense data’ to which one can apply the terms ‘iron’ and ‘tower’ via more complex logical statements involving towers and their descriptions. Likewise, the idea of direct acquaintance with sense data equally underpins both Putnam and Berners-Lee, as both think that reference should be determined by expertise, for instead of just labeling a patch of sense-data with the term ‘iron tower,’ the scientists would label the sense-data with the use of a name like ‘iron tower’ only after it successfully passed some authoritative test, such as a test for the chemical composition of iron.

Using the example of the ‘duck-rabbit’, Wittgenstein undermines the very idea of establishing a referent via direct acquaintance and baptism (Wittgenstein, 1953). If one can not determine that a simple sketch is of a ‘duck’ or a ‘rabbit,’ then how can *anyone* objectively and without ambiguity attach a name to some sense-data? The indeterminacy of the infamous ‘duck-rabbit’ shows that at least in some cases there is no determinate nature of our phenomenological ‘sense-data.’ Having disposed of the notion of ostention somehow providing direct access to sense-data, baptism of even indeterminate sense-data – by either Kripkean baptism or Russellian descriptions – is attacked next. Wittgenstein holds that any act of baptism is incapable of assigning a name if the act is done by a private individual, “naming appears as a *queer* connection of a world with an object – and you really get such a queer connection of a word when a philosopher tries to bring out the relations between name and thing by staring at an object in front of him and repeating a name or even the word ‘this’ innumerable times” (Wittgenstein, 1953). Only in the very rarefied situations does this happen because “naming is so far not a move in the language-game any more than putting a piece in its place on a board is a move in chess. We may say: *nothing* has so far been done, when a thing has been named. It has not even *got* a name except in the language game. This is what Frege meant too, when he said that a word has meaning only as part of a sentence” (Wittgenstein, 1953). Indeed, naming as a purely private convention serves no purpose. It is only as part of a wider language that anything can have a name in the first place. Even what appears to be the most private of sensory experiences is both determined and expressed by a public language.

As championed by Searle’s account of “social reality” in natural language, a name has meaning via

the use and acceptance of speech acts in public language, so that a Kripkean ‘baptism’ is only one kind of a wide array of possible speech acts. To broaden our horizon, descriptions can be considered another kind of speech act in public language Searle (1995). Without this acknowledgment of a public language creating a shared objective social reality via its effects and usage by participating agents, we have no choice but to ascribe certain seemingly magical referential powers to ‘baptism’ or mysteriously connecting logical descriptions with sense data via direct acquaintance. The magical and mysterious connection between words and meaning is revealed by the public language position to be nothing other than the deployment of language convention and innovation.

In contrast to Wilk’s interpretation of Wittgenstein and computational linguistics Wilks (2008b), we do not hold that the “magic” of meaning be accounted for by some ‘mental’ lexicon or even massive statistics on word-use on the Web. Taking the role of embodiment and the world seriously, it would be a mistake to separate language from the world and the actions of agents therein. Therefore, natural language processing itself is not a firm foundation for semantics on the Web. On the contrary, it is of crucial to attach the use of words and meaning to particular embodied language-games, like the use of search engines. After all, to use words to retrieve web-pages is vastly different than using words in dialogue with another human. Many of the characteristics of terms typed into search engines, such as their extreme brevity, will have no clear analogue to natural language use. The open-ended innovative use of language on the Web makes the creation of finite lexicons difficult if not impossible.

7.2 Public Language and Ambiguity

The question that bedeviled the causal theory of reference was that of ambiguity. Does ambiguity result from a failed use of language? Ambiguity is built into a public language position, and the kind of ambiguity that Wittgenstein is concerned with is not the kind of ambiguity resulting from entailments failing to constrain interpretations. For Wittgenstein, ambiguity is naturally constrained by the conventions of the language, which are restricted in turn by the external world. While the Wittgensteinian public language position would note that there is always some ambiguity in language, worrying about this ambiguity misses the point, as the point of a language game is not to pin down names to referents exactly, but instead to share enough of a convention to accomplish some task or solve some problem. Ambiguity is usually solved by the embodied or

implicit context given in the language – it is not without reason that Wittgenstein begins the *Philosophical Investigations* contrasting the Augustinian approach of assigning the builders to objects with the “language game” of builders moving slabs or rock around. For the builders, their task at hand determines their meaning of the word. The role of descriptions and inference is not in determining referents, but only when the various agents in a language-game are not clear about the role of a name in a language game, so that “an explanation may indeed rest on another one that has been given, but none in need of another – unless *we* require it to prevent a misunderstanding” (Wittgenstein, 1953). In this manner, inference and entailments that restrict interpretations, as defended by Hayes, are a logical analogue to the real-world context that both constrains ambiguity in a language game, while usually never dispelling it. While some inferential mechanisms can be useful when errors are made in a language game, in general inference can not express all the constraints given by the contextual use of a name in a language game. However, it is a beginning, and an important one.

How does the public language position actually play out on the Semantic Web? To apply Wittgenstein to the Semantic Web, the first observation is that the Semantic Web *is* a *new* language-game. There is no reason why language-games in a Wittgensteinian sense have to be restricted to natural languages, for Wittgenstein himself notes that “new types of language, new language-games, as we may say, come into existence, and others become obsolete and get forgotten” (Wittgenstein, 1953). The struggle over the Identity Crisis within the Semantic Web is precisely the struggle over the conventions of meaning needed for a new language. We suppose that the terms ‘language’ and ‘sense’ can be *neutral* between formal languages for computers and natural languages. Formal languages are often mistakenly thought to be meaningless due to their not taking into account the concrete activity that occurs as a result of their use but instead to be pure “syntax churning” (Harnad, 1990). Given that agents can be computers as much as humans, with their own norms for behavior – such as protocols – there seems to be no reason why computers, or combinations of computers and humans, cannot create and use new language-games. So from the perspective of the public language position, when a new URI comes into play on the Semantic Web, the agents do not have to specify the referents of the URI to use it meaningfully. If the referent of a name has to be specified for the name to be used, it only has to be specified to the minimal conditions necessary to co-ordinate actions between agents. Contra Berners-Lee’s direct reference position, only in very rare language games does

the referent of some representation have to be specified in an ‘unambiguous’ manner.

7.3 Search Engines

While we could sketch how a theory of meaning based on the public language position could work for the Semantic Web, are there any real examples of how such a theory of meaning actually *does* work? The answer, rather surprisingly, comes from existing Web search engines like Google. Information retrieval, and its data-driven methodology, *are* neo-Wittgensteinian philosophy of language given computational flesh, as held by Wilks Wilks (2008b). The discipline of information retrieval is directly descended from Wittgenstein himself via the under-appreciated philosopher and linguist Margaret Masterman. The history of how Wittgenstein influenced Web search engines is a fascinating trajectory. Wittgenstein’s infamous dictum that “meaning is use” seems often itself meaningless upon first glance; how can “meaning is use” possibly be operationalized into a methodology that could form the basis for a science of language (Wittgenstein, 1953)? The answer is obvious: in studying the structure of language empirically, which can be done computationally by the statistical analysis of actual samples of human language. In other words, the building of “language processing programs which had a sound philosophical basis” (Wilks, 2005). One of the six students of Wittgenstein’s course that became *The Blue Book*, Masterman was exposed directly by Wittgenstein to the conceptual apparatus of the *Philosophical Investigations* (Sowa, 2006). Twenty years later, she founded the Cambridge Language Research Unit, where the foundations for information retrieval were laid by a student of Masterman and her Richard Braithwaite, Karen Spärck Jones (Wilks, 2007). In her dissertation *Synonymy and Semantic Classification*, Spärck Jones stated that her dissertation proposed “a characterisation of, and a basis for deriving, semantic primitives, i.e. the general concepts under which natural language words and messages are categorized.” (Jones, 1964). She did this by applying the statistical ‘Theory of Clumps’ of Roger Needham – a theory that was itself one of the first to explicate what Wittgenstein called “family resemblances” – to words themselves, leading her to posit that words could be defined in terms of statistical clumps of other words (Needham, 1962). Applying her work over larger and larger sources of natural language data, she later began to abandon using even the open-ended semantic primitives of Masterman. In her later critique of artificial intelligence, she cited that one of the key insights of information retrieval is that programs should take “words as they stand” and

not as only adjuncts to some logical knowledge representation system (Jones, 1999). Altavista, the first of modern Web search engines, was created after its inventor, Mike Burrows, e-mailed Spärck Jones and Needham over techniques in information retrieval.

Search engines work via analysis of existing web-pages, breaking them down into terms, and then mapping those terms and their frequencies in a given web-page into a large index. So, each URI can be thought of as collection of terms in this search engine index. As the collection of term frequencies gathered into this index grows, ranging over larger and larger sources of data like the Web, it approximates a sample of human language use, as has been shown by studies in computational linguistics (Keller and Lapata, 2003). Users of a search engine then enter certain terms, the search query, which are then mapped via certain algorithms against the index. This results in an unordered list of possibly relevant URIs, which for an index that covers the entire Web range from thousands to millions of URIs. In turn these URIs are then ranked and ordered using an algorithm such as Google’s famous PageRanking algorithm, possibly with user feedback (Brin and Page, 1998). To explicate how user feedback works, search engines usually keep track on what URIs are actually clicked on by users. This stream of clicks by multiple users can then be stored in a “query log,” and then this query log can then be used to improve the discovery and ranking of URIs by search engines. By inspecting which terms lead to which URIs for multiple users, a set of terms that best describes a URI for users can be discovered. In this way, typing in terms into a search engine can be thought of as a sort of language-game, with success in this game being judged in terms of whether or not a given user can, using a particular Web search engine, discover a relevant URI. The terms themselves may be ambiguous, but it does not matter if a relevant URI is discovered. While the user may not be aware of it, as it *appears* that searching the Web using a search engine is a private experience, it is in fact mediated by a vast amount of web-pages stored in the search engine’s index and the behavior of previous search users. In this way, the objective sense of a URI can be considered the search terms that can be used by multiple users to find a particular URI. What a URI means is precisely the set of search terms that leads multiple users to discover the URI in the context of satisfying a particular information need. It should also be noted that this position does not overtly contradict the logical position, as the logical position radically undermines the referent(s) on the Web. Instead, one can imagine that the public language position allows a URI to be grounded in user-

behavior using search terms, but that this meaning can be supplemented by logical inference.

8 Conclusion: A Statistical Semantic Web

The Semantic Web has yet to be widely deployed, and it could be precisely due to the persistence of its the problems of the philosophy of language regarding meaning and reference that it attempts to build upon. We have argued that the debates over the meaning and reference of URIs can be seen as a return of the debate between the causal and descriptivist theories of reference in the philosophy language, with this time the subject being URIs rather than natural language names. In this way, the it has been shown that in the course of the practice of computer science, even in such a new undertheorized and undisciplined frontier like the Web, robustly philosophical problems arise. By stumbling on the difficult philosophical problem of reference and meaning, it appears that the success of the Semantic Web, one of the most ambitious projects of knowledge representation so far, has been stymied. Indeed, the deadlock between Hayes and Berners-Lee on the meaning of URIs reflects a long-standing debate within philosophy.

What does the practice of computer science on the Web have to offer to these debates in the philosophy of language and on the Web? The observation from this work is that the infamously cryptic ideas of sense and Wittgenstein's maxim "meaning is use" can be grounded in computation, and this methodology may provide a way to go beyond deadlocks of theories of reference and naming on the Web. We have argued that Web search engines are the direct descendants of a competing Wittgensteinian theory of language, on that tries to leverage the meaning of URIs not in their attachment to either logical descriptions or referents, but in their sense as given by the terms used to discover these URIs. Given the massive success of search engines like Google in allowing users to structure the Web, it may be that it is the time to revive the philosophical debates around meaning and reference in a more computationally informed manner, with attention being paid to how the computational embodiment of these theories have played out "in the wild." The central position of hypertext search engines should not be underestimated on the Web, and we find it astounding that the Semantic Web has ignored hypertext Web search engines up till very recently. While the techniques behind information retrieval and machine-learning may not be cognitively transparent, it is perhaps an error of both the descriptivist and causal theories of reference to hold that meaning can be reduced to straightforward accounts. In fact,

acknowledging the kinds of complex social-technical interactions exemplified by search engines may indeed be more likely to do justice to the problem of attaching meaning to words.

Searching the large hypertext Web leads precisely to a 'statistical semantic web,' where the meaning of URIs is given by the activity of users. In this way, the bet of using URIs as a universal naming scheme for things can just as easily be tied to statistical methods from information retrieval as it can to logic-based knowledge representations. In this way, the philosophy of computer science can even provoke further practical engineering work on the Web, and it is precisely the success of computational engineering that wins debates on the Web, rather than pure argumentation. Currently information retrieval engines like Google far outweigh the Semantic Web in terms of usage, which give us a clue as regards which philosophical position may be correct. Is there any way the massive amount of search terms and URIs harvested by search terms be used to boot-strap the Semantic Web? One bet would be that when users want to actually find URIs data on the Semantic Web or even find URIs to re-use for labeling their data on the Semantic Web, they will have to employ search-engines over the Semantic Web with the same natural language-driven searches currently employed by Web search engines over the hypertext Web. However, instead of hypertext web-pages being indexed, the index will range over decentralized knowledge representation stored in RDF. A combination of techniques of search engines that approximate the meaning of URI using natural language terms derived from the use of natural language text in the knowledge representation itself (ideally, with some machine translation) with the inference-driven capabilities of RDF would be ideal. This could lead to a new kind of searching, not just over web-pages, but over actual facts about real-world things, 'searching for things' rather than 'searching for hypertext.' Some aspects of this can be seen in recent work on search engines by Wolfram and the focus of even Google on 'semantic search,' where Web search goes beyond just natural language terms. Indeed, it is precisely in this direction that much recent work on the Semantic Web has been aimed in grounding URIs on the Semantic Web in their actual discovery and use (Oren et al., 2008). Unfortunately due to the hold of the descriptivist and causal theories of reference on the minds behind the Semantic Web, the 'Semantic' Web has no meaning for most users, but is a meaningless jumble of URIs threaded to together by a hard-to-understand knowledge representation system. As information retrieval is only one kind of interaction with the Web, attaching the Semantic Web to more rich Web-mediated social

interaction is exactly what is necessary to add meaning to any knowledge representation.

If there is anything to be learned from this encounter of philosophy and the Web, it is that one can never escape philosophical problems, even on the Web. They cannot simply be ignored. The philosophy of language has had a large influence on knowledge representation languages in general, in particular the influence of logical theories of reference on artificial intelligence that has continued to influence the Web via Pat Hayes. Yet strangely enough, it is the philosophically untrained Berners-Lee that ends up arguing what has been accepted in philosophical circles as the causal theory of reference, a theory of reference that is widely accepted in some circles to be correct. Even more surprisingly, it appears that search engines like Google embody an alternative theory of meaning, one based on an objective notion of sense implicitly given by Wittgenstein. Indeed, what is necessary is not only an illumination of debates on the Web in terms of philosophy, but the revival of these debates within the framework of the philosophy of computer science. The exact relationship between the philosophy of computer science and the philosophy of language is a vast terrain that remains to be explored. The results of such an exploration are not just of philosophical interest, but of interest to the practice of computer science on the Web.

Acknowledgements I would like to thank the many fruitful conversations I've had on these topics with Tim Berners-Lee and Pat Hayes that revealed the fundamentally philosophical nature of their disagreement, as well as encouragement and comments when this work has been presented in public, both at the 'Identity, Reference, and the Web' workshop at the World Wide Web Conference and at the North American Philosophy and Computing Conference. In particular, substantial comments from Henry S. Thompson and Yorick Wilks have been of worked into the text.

References

- Beckett, D. and Berners-Lee, T. (2008). Turtle - Terse RDF Triple Language. Member submission, W3C.
- Berners-Lee, T. (1994). World Wide Web Future Directions. Plenary Talk. <http://www.w3.org/Talks/WWW94Tim/> (Last accessed on Oct. 5th 2008).
- Berners-Lee, T. (1998). Cool URIs don't Change. <http://www.w3.org/Provider/Style/URI> (Last accessed on Nov 19th 2008).
- Berners-Lee, T. (2003a). Message on www-tag@w3.org list. <http://lists.w3.org/Archives/Public/www-tag/2003Jul/0158.html> (Last accessed on May 20th 2008).
- Berners-Lee, T. (2003b). Message to www-tag@w3.org. <http://lists.w3.org/Archives/Public/www-tag/2003Jul/0022.html> (Last accessed on May 20th 2008).
- Berners-Lee, T., Fielding, R., and Masinter, L. (1998). IETF RFC 2396 Uniform Resource Identifier (URI): Generic Syntax. <http://www.ietf.org/rfc/rfc2396.txt> (Last accessed on Sept. 15th 2008).
- Berners-Lee, T., Fielding, R., and Masinter, L. (2005). IETF RFC 3986 Uniform Resource Identifier (URI): Generic Syntax. <http://www.ietf.org/rfc/rfc3986.txt> (Last accessed on April 2th 2008).
- Berners-Lee, T., Hendler, J., and Lassila, O. (2001). The Semantic Web. *Scientific American*, 284(5):35–43.
- Brin, S. and Page, L. (1998). The anatomy of a large-scale hypertextual Web search engine. In *Proceedings of the International Conference on World Wide Web (WWW)*, pages 107–117, Brisbane, Australia.
- Carnap, R. (1928). *The Logical Structure of the World*. University of California Press, Berkeley, California, USA. Republished in 1967.
- Dummett, M. (1993). What is a Theory of Meaning. In *The Seas of Language*, pages 1–33. Oxford University Press, Oxford, United Kingdom. Originally published in *Truth and Meaning: Essays in Semantics* in 1976.
- Frege, G. (1892). Uber Sinn und Bedeutung. *Zeitschrift fur Philosophie und philosophie Kritik*, 100:25–50. Reprinted in *The Philosophical Writings of Gottlieb Frege* (1956), Blackwell, Oxford, United Kingdom (1956), translated by Max Black.
- Grice, P. (1957). Meaning. *The Philosophical Review*, 66:377–388.
- Guha, R. V. (1996). Meta Content Framework: A White paper. <http://www.guha.com/mcf/wp.html> (Last accessed Aug. 11th 2008).
- Halpin, H. and Lavrenko, V. (2009). Relevance feedback between hypertext search and semantic search. In *Proceedings of the Semantic Search Workshop at the World Wide Web Conference*. Madrid, Spain.
- Halpin, H. and Thompson, H. (2005). Web Proper Names: Naming Referents on the Web. In *Proceedings of The Semantic Computing Initiative Workshop at the World Wide Web Conference*, Chiba, Japan.
- Harnad, S. (1990). The Symbol Grounding Problem. *Physica D*, 42:335–346.
- Hayes, P. (1977). In defense of logic. In *Proceedings of International Joint Conference on Artificial Intelligence*, pages 559–565, Cambridge, Massachusetts, USA.
- Hayes, P. (2003a). Message to www-rdf-comments@w3.org. <http://lists.w3.org/Archives/Public/www->

- tag/2003Jul/0147.html (Last accessed on May 20th 2008).
- Hayes, P. (2003b). Message to www-rdf-comments@w3.org. <http://lists.w3.org/Archives/Public/www-tag/2003Jul/0198.html> (Last accessed on May 20th 2008).
- Hayes, P. (2004). RDF Semantics. Recommendation, W3C. <http://www.w3.org/TR/rdf-mt/> (Last accessed Sept. 21st 2008).
- Hayes, P. (2006). In defense of ambiguity. In *Proceedings of the Identity, Reference, and the Web Workshop at the WWW Conference*, Edinburgh, Scotland. <http://www.ibiblio.org/hhalpin/irw2006/hayes.pdf> (Last accessed on Oct. 5th 2008).
- Hayes, P. and Halpin, H. (2008). In defense of ambiguity. *International Journal of Semantic Web and Information Systems*, 4(3).
- Jones, K. S. (1964). Synonymy and semantic classification. Thesis, Cambridge University. Republished in 1984 by Edinburgh University Press.
- Jones, K. S. (1999). Information Retrieval and Artificial Intelligence. *Artificial Intelligence Journal*, 114:257–281.
- Keller, F. and Lapata, M. (2003). Using the web to obtain frequencies for unseen bigrams. *Computational Linguistics*, 29(3):459–484.
- Kripke, S. (1972). *Naming and Necessity*. Harvard University Press, Cambridge, Massachusetts, USA.
- Luntley, M. (1999). *Contemporary Philosophy of Thought*. Blackwell, London, United Kingdom.
- Needham, R. (1962). A method for using computers in information classification. In *Proceedings of the IFIP Congress*, pages 284–287, Vienna, Austria.
- Oren, E., Delbru, R., Catasta, M., Cyganiak, R., Stenzhorn, H., and Tummarello, G. (2008). Sindice.com: A document-oriented lookup index for open linked data. *International Journal of Metadata, Semantics, and Ontologies*, 3(1):37–52.
- Putnam, H. (1975). The meaning of meaning. In Gunderson, K., editor, *Language, Mind, and Knowledge*. University of Minnesota Press, Minneapolis, Minnesota, USA.
- Russell, B. (1905). On denoting. *Mind*, 14:479–493.
- R.V.Guha and D.Lenat (1993). Language, representation and contexts. *Journal of Information Processing*, 15(3).
- Searle, J. (1995). *The Construction of Social Reality*. The Free Press, New York City, New York, USA.
- Smith, B. C. (1995). *The Origin of Objects*. MIT Press, Cambridge, Massachusetts, USA.
- Sowa, J. (2006). Review of language, cohesion, and form by margaret masterman. *Computational Linguistics*, 4(32):551–553.
- Tarski, A. (1935). The concept of truth in formalized languages. *Studia Philosophia*, 1:261–405. Reprinted in *Logic, Semantics and Metamathematics* (1956), Oxford University Press, Oxford United Kingdom, (1956), translated by J.H. Woodger.
- Tarski, A. (1944). The semantic conception of truth and the foundations of semantics. *Philosophy and Phenomenological Research*, 4:341–375.
- Wilks, Y. (2005). A personal memoir: Margaret Masterman (1910-1986). In Masterman, M., editor, *Language Cohesion and Form*. Cambridge University Press, Cambridge, United Kingdom.
- Wilks, Y. (2007). Karen Spärck Jones (1935-2007). *IEEE Intelligent Systems*, 22(3):8–9.
- Wilks, Y. (2008a). The Semantic Web: Apotheosis of annotation, but what are its semantics? *IEEE Intelligent Systems*, 23(3):41–49.
- Wilks, Y. (2008b). What would a Wittgensteinian computational linguistics be like? In *Proceedings of Convention for the Society for the Study of Artificial Intelligence and Simulation of Behavior*, Aberdeen, Scotland.
- Wittgenstein, L. (1953). *Philosophical Investigations*. Blackwell Publishers, London, United Kingdom. Republished 2001, translated by G.E.M. Anscombe.
- Woods, W. (1975). What’s in a link: Foundations for semantic networks. In *Representation and Understanding: Studies in Cognitive Science*, pages 35–82. Academic Press, Inc., Orlando, Florida, USA.