Tracing the origins of the semantic web

Reference:
Guns Raf.- *Tracing the origins of the semantic web*
Journal of the American Society for Information Science and Technology / American Society for Information Science and Technology - ISSN 1532-2882 - 64:10(2013), p. 2173-2181
DOI: http://dx.doi.org/doi:10.1002/asi.22907
Handle: http://hdl.handle.net/10067/1111170151162165141

This paper is a postprint of the following article:

Guns, R. (2013). Tracing the origins of the Semantic Web. *Journal of the American Society for Information Science and Technology*, 64(10), 2173–2181.
http://dx.doi.org/10.1002/asi.22907

Tracing the origins of the Semantic Web

Raf Guns

Universiteit Antwerpen, IBW, Venusstraat 35, B-2000 Antwerp, Belgium

Abstract

The Semantic Web has been criticized for not being semantic. This article examines the question: why and how has the Web of Data, expressed in RDF, come to be known as the *Semantic* Web? Contrary to previous papers, we deliberately take a descriptive stance and do not start from preconceived ideas about the nature of semantics. Instead, we mainly base our analysis on early design documents of the (Semantic) Web.

The main determining factor is shown to be link typing, coupled with the influence of online metadata. Both factors are already present in early Web standards and drafts. Our findings indicate that the Semantic Web is directly linked to older AI work, despite occasional claims to the contrary.

Because of link typing, the Semantic Web can be considered an example of a semantic network. Originally network representations of the meaning of natural language utterances, semantic networks have eventually come to refer to any networks with typed (usually directed) links. We discuss possible causes for this shift and suggest that it may be due to confounding paradigmatic and syntagmatic semantic relations.

Introduction

In recent years, the Semantic Web has seen considerable attention in a wide variety of research fields, including information science (Legg, 2007; O'Hara & Hall, 2010; Rousseau & Rousseau, 2002). The Semantic Web is essentially a new development in the realm of representation and organization of knowledge, one of the core subfields of information science. Hence, the Semantic Web and associated standards form a very relevant study object of information science. At the same time, the Semantic Web has drawn heavily on earlier insights and techniques from information science and related fields like librarianship: technologies and standards such as SKOS or Dublin Core have roots in library and information science.

Despite early enthusiasm, the Semantic Web has been subject to criticism as well (e.g., Brooks, 2002; Shirky, 2003). Some of the harshest criticisms have been leveled against the *name* ‘Semantic Web’ or, perhaps more accurately, against the Semantic Web’s apparent claim of involving semantics. Indeed, it is not intuitively clear in what way the Semantic Web can be said to be ‘semantic’: over the past decade, several authors have criticized the Semantic Web and its proponents for disregarding or misrepresenting the nature of its semantics. The word ‘semantic’, deriving from Greek *semantikós*, signifies ‘pertaining to meaning’. Semantics is the study of meaning, especially the relation between signifier (e.g., a word) and signified (e.g., what a word refers to in reality). Semantics forms one of the three main ways in which a language can be studied, the other two being syntax (the study of how units can be grammatically related) and pragmatics (the study of how units relate to the context in which they are used). Stated another way, semantics can be characterized as the study of how (linguistic) units relate to items in the world. Semantics is studied in linguistics and psychology, as well as in fields like logic, artificial intelligence (AI) and computer science. In the latter cases, it may relate to natural languages (e.g., Natural Language Processing) or formal languages (e.g., programming languages). These examples illustrate that semantics cuts across a broad range of fields. Hence, caution should be exercised when one encounters an unfamiliar manifestation of the word ‘semantic’, in that it may not always refer to the same (or even similar) concepts.

Gärdenfors (2004) claims that “the Semantic Web is not semantic,” mainly on account of its limitations to first order logic or set theory: “If one considers how humans handle concepts, the class relation structures of the Semantic Web capture only a minute part of our information about concepts. For instance, we often categorize objects according to the *similarity* between the objects […]. And similarity is not a notion that can be expressed in a natural way in a web ontology language.” (Gärdenfors, 2004, p. 18–19). Uschold (2003) notes that many Web applications work without any significant regard for semantics, citing the example of shopping agents. Such applications work because human programmers have hardwired the form and meaning of a specific type of Web content into the application. He suggests that the way forward is a move towards more explicit, formal semantics.

Harnad (2002) argues that the Semantic Web should be called the 'Syntactic Web' because it suffers from the *symbol grounding problem*: a Semantic Web symbol (e.g., a URI) is not inherently connected to the object to which it refers or, in other words, its meaning is derived from an external, human interpreter. Sparck-Jones (2004) believes that, in order to fulfill the Semantic Web vision (Berners-Lee, Hendler, & Lassila, 2001), one would need a high-end, all-encompassing world model and that it would be better "to bite the natural language bullet right away."

In our opinion, the discussion whether or not the Semantic Web 'is really semantic' is unfruitful, because it is dependent on each commenter's own biases and opinions regarding what semantics is and/or should be. Consider, for instance, the question in what direction the Semantic Web should evolve. While Gärdenfors (2004) presents arguments for moving away from formal semantics, Uschold's (2003) opinion is that *more* formal semantics are needed. In other words, these two authors are almost diametrically opposed on this matter, even though at first glance they might seem to agree and, ironically, the former article cites the latter.

In this article, we take a more descriptive route, trying to answer the question where the notion of semantics in relation to the Semantic Web comes from, rather than discussing whether or not this is appropriate usage. More specifically, we trace the Semantic Web's origins by considering early design documents. It seems that the notion of semantics does not primarily derive from ontologies as such (even though this is often thought to be the case), but is mainly related to link typing as in *semantic networks*.

Almeida, Souza, and Fonseca (2011) consider the question "what kind of semantics we can find in the Semantic Web." They discuss existing theories of semantics and characterize Semantic Web technologies along a spectrum of semantic expressiveness. Although the question of the current paper is in fact very similar, our approach is rather different. Specifically, we rely almost exclusively on documents that originate within the Semantic Web community and do not try to shoehorn Semantic Web technologies into pre-existing theories of semantics.

The remainder of the article is structured as follows. In the next section, we discuss the historical background of the Semantic Web, focusing in succession on original Web proposals, semantics in HTML, and the connection with CycL via MCF. The following section examines semantic networks in more detail. Subsequently, we discuss the implications for the semantics of the Semantic Web. The final section presents the conclusions.

Historical background of the Semantic Web

The World Wide Web is the largest information system in the history of mankind. It was developed in 1990 at the CERN research facilities in Genève, Switzerland by Tim Berners-Lee and Robert Cailliau to improve the sharing of research data and results between scientists (Berners-Lee, 2000). Although the Web is mainly known for documents in HTML, even the earliest proposals hint at ambitions beyond a simple document formatting language: “The vision of a web of data was always implicit in the ideas underlying the development of the WWW” (O’Hara & Hall, 2010, p. 4667). The Semantic Web does therefore not represent a radical paradigm shift, but rather the externalization of an undercurrent that has existed since the beginning of the Web. Throughout the first decade of the Web, several proposals have been developed to enable ‘semantic’ relationships on the Web. In this section, we provide an overview of the most important ones, and discuss how the current design has developed.

Ideas in original Web proposals

Close examination of the original World Wide Web proposals reveals that these documents already contain many of the ideas underlying the Semantic Web. In his earliest proposal to introduce an Internet-based hypertext system (Berners-Lee, 1989), Tim Berners-Lee describes the well-known problems of scattered information, incompatible file formats and changing information needs. He proposes that a hypertextual information system may resolve these problems. This diagram still uses the term ‘Mesh’, which would later be replaced by ‘World Wide Web’.

The Web is generally conceived as a network of interlinked documents (Björneborn & Ingwersen, 2001; Thelwall, 2004; Thelwall, Vaughan, & Björneborn , 2006). It is therefore rather surprising that the figures in (Berners-Lee, 1989) depict connections between documents, concepts, computer systems, people, and organizations. Berners-Lee describes the system as a network with nodes and links, where nodes and links “can stand for anything.”

His examples of possible nodes include: people, software modules, projects, concepts and documents. Links as well can be of different types; a link from A to B “can mean, for example, that A depends on B; is part of B; made B; refers to B; uses B; is an example of B.” Further on, he remarks: “In practice, it is useful for the system to be aware of the generic types of the links between items (dependences, for example), and the types of nodes (people, things, documents..) without imposing any limitations” (Berners-Lee, 1989).

In summary, key ingredients of the proposed system are:

- the system is agnostic regarding its content – that is, any kind of information can be entered;
- it exhibits a network-like structure;
- one node represents one object;
- both links and nodes are of different types.

Some of the early design documents from 1990 for the Web explicitly discuss the question whether links should have types and what types should be available. It is remarkable that these early documents explicitly connect typed links to semantics: “A typed link carries some semantic information, which allows the system to manage data more efficiently on behalf of the user.”[1]

The same connection has been carried through in early design documents for the Semantic Web: typed links are considered to carry semantics. A good illustration of this connection was found when trying to answer the question: when was the Semantic Web first mentioned? Note that the *term* is certainly older than the *concept* as it is used here. Indeed, a simple search of some scientific literature databases reveals that the term ‘semantic web’ was used already in fields like educational psychology. Freedman and Reynolds (1980), for instance, refer to ‘semantic webbing.’

If we restrict ourselves to mentions of the Semantic Web in the intended sense, the oldest available document we could find dates back to June 1995. In a keynote speech at INET’95, the Internet Society’s 1995 International Networking Conference, Tim Berners-Lee outlined four points on the way to a ‘semantic web’.[2] The full text of the presentation is not available, but the four points mentioned in the surviving slides are:

1. Link typing

2. Knowledge representation content types
3. Meta language for trait investigation
4. Bootstrapping class structures

It is not entirely clear what exactly was meant by each of these points, but the first one is quite clear: typed links imply semantics.

Remarkably, this implies that the first public mention of the Semantic Web happened only four years after the birth of the Web itself and six years before the well-known *Scientific American* article that made the project known to a much wider audience (Berners-Lee et al., 2001). In fact, the same points were used even earlier by Berners-Lee in a presentation from 1994 for WWW Fall 94, but there they were mentioned under the heading 'Semantics' instead of 'Semantic Web'.

Metadata and 'semantics' in HTML

Another factor that has influenced the Semantic Web is metadata on the Web. An orientation towards metadata has always been part of the Web. There has been (and still exists to this very day) a strong push towards encoding metadata into HTML, the language of Web documents. This push was both internal, championed by Berners-Lee and others at the World Wide Web Consortium, and external, by many different actors. Some of these developments are also discussed in an early overview paper by Alexander Chislenko (Chislenko, 1997). Interestingly, Chislenko arrived at the same terminology ('Semantic Web') apparently independently from Berners-Lee's work. His paper touches upon many themes, but most striking is his description of 'relations between items', where he mentions relations between documents, but also between people and between people and companies, – some years later the FOAF project made this vision come true. Chislenko also mentions the importance of item ontologies and stresses the importance of 'compatibility of multiple ontologies'. This is, in fact, still a major issue nowadays.

The Hypertext markup Language (HTML) was initially considered a relatively unimportant part of the young Web (Berners-Lee, 2000). Consequently, HTML did not have a very elaborate or strict technological specification. This changed with the introduction of HTML 2.0 in 1995. While fairly primitive compared to current iterations of the language, HTML 2.0 (Berners-Lee & Connolly, 1995) already contains the seeds of metadata embedding and typed linking.

Metadata can be added in HTML 2.0 using the aptly labeled `meta` element, which has two goals: “to provide a means to discover that the data set exists and how it might be obtained or accessed” and “to document the content, quality, and features of a data set, indicating its fitness for use.” For example, the author of a Web document could be specified as follows:

```
<meta name="author" content="John Doe">
```

Thus, the `meta` element, which still exists, provides for a simple, crude but efficient way of adding key–value metadata to Web pages. An obvious downside is that the value of the `name` attribute is open to interpretation: does my `name="type"` signify the same as yours? Later specifications have introduced the concept of a *profile*: a reference to an external document that specifies what properties exist, what their values can be, and so on. The latest iteration of HTML, HTML 5 (Hickson, 2012a; currently in draft), abandons profiles, but provides other options to incorporate metadata into a Web document.

Typed linking is possible in HTML 2.0 with the `rel` and `rev` attributes: “The `REL` attribute gives the relationship(s) described by the hyperlink. The value is a whitespace separated list of relationship names. The semantics of link relationships are not specified in this document.” The `rev` attribute also describes the link relationship but in the opposite direction: “[a] link from A to B with `REL="X"` expresses the same relationship as a link from B to A with `REV="X"`.” HTML 5 (Hickson, 2012a) has obsoleted the `rev` attribute, because it was deemed confusing and turned out to be hardly used in practice.[3] `rel` is used more often, although mainly for a limited set of specific technical use cases, such as CSS stylesheets or RSS/Atom feeds.

Features like the `meta` element and `rel` and `rev` attributes provide basic possibilities to add metadata to Web pages and enhance interpretation by software. Of course, these features are also rather limited. For instance, a fairly common need that these features do not address is indicating that a specific part of a Web page is a person’s name, a date, an address etc. To a certain extent, this need can be alleviated by introducing specialized HTML elements, – HTML 5 introduces a `time` element, for example. However, each new element increases the complexity of the language and it is impossible to introduce new elements for *any* need.

Several solutions have been considered for this problem. Perhaps the first genuine attempt to add a ‘semantic’ component to HTML is SHOE (Simple HTML Ontological Extensions), described in (Heflin & Hendler, 2000; Heflin, 2001; Luke, Spector, & Rager, 1996) and on http://www.cs.umd.edu/projects/plus/SHOE/. SHOE is essentially an extension of HTML that allows specifying data objects, their classes, and their relations. Consider the following SHOE fragment:

```
<INSTANCE KEY="http://www.example.com/people/#john">
    <USE-ONTOLOGY
     URL="http://www.example.org/ontology/people.html"
     PREFIX="pp">
    <CATEGORY NAME="pp.Farmer">
    <RELATION NAME="pp.name">
        <ARG POS=TO VALUE="John Doe">
    </RELATION>
    <RELATION NAME="pp.age">
        <ARG POS=TO VALUE="42">
    </RELATION>
</INSTANCE>
```

Here, we define an object (an INSTANCE) which is assigned a Web address (a URI, see Horrocks, Patel-Schneider, & van Harmelen, 2003). To describe it, we import an ontology that exists elsewhere on the Web. We then declare that the object belongs to a category pp.Farmer and create two relations (name and age). Note that the prefix pp is also associated with a URI.

Around the same time, Kogan, Michaeli, Sagiv, and Shmueli (1998) propose that semantic elements can be added in HTML comments: thus, they are invisible to normal browsers but can be processed with specialized software. Their proposal, named OHTML, is noteworthy for recognizing (a) that all objects need a unique object identifier, and (b) that a multi-relational network functions well as the underlying data structure. Nevertheless, it has never gained much traction, presumably because of the complexity of adding structured data to HTML comments.

More recent solutions to the problem of how to make parts of a Web page machine-processable include *microformats* (http://microformats.org; see Khare & Çelik, 2006), *RDFa* (http://rdfa.info/), and *microdata* (Hickson, 2012b). Work is currently ongoing to unify these complementary approaches.

CycL, MCF, PICS and RDF

We have discussed early initiatives involving link typing and metadata on the Web. However, it would be an oversimplification to say that RDF, the language of the Semantic Web, is directly derived from these initiatives. In fact, RDF's design is heavily influenced by work that is rooted in traditional AI: MCF and, indirectly, CycL. We now focus on this connection.

One of the most ambitious ontological engineering efforts ever undertaken, *Cyc* is a large-scale 'common-sense' ontology and knowledge base (Lenat, 1995); the name derives from 'encyclopedia'. Cyc's primary author is Douglas Lenat, who has been developing the system since 1984 up until today (along with many others). In 1994, Lenat founded Cycorp, Inc. as a commercial successor to the original Cyc Project. OpenCyc[4] is an open source version of the Cyc knowledge base.

Knowledge in Cyc is stored as assertions in the CycL representation language (Lenat & Guha, 1991). Currently, the system contains about five million assertions.[5] These are typically of the form 'Apples are edible fruit' or 'Edible fruit is food'. The combination of these two assertions would lead the system to deduce that apples are food and can be eaten. In CycL, they would be written as:

```
(#$genls #$Apple #$EdibleFruit)
(#$genls #$EdibleFruit #$Food)
```

`#$genls` is one of the most important predicates in CycL, expressing the generic (genus–species) relation. That is, it states that the first argument (e.g., `#$Apple`) is a subclass of the second one (e.g., `#$EdibleFruit`). It is the CyCL equivalent of the hypernym–hyponym relation. Assertions in CycL are *n*-tuples, where usually $n = 3$ (triples). The tuple structure and syntax derive from so-called S-expressions in the Lisp programming language, which was – and still is – very popular for AI research. Each assertion is associated with a microtheory (a context) and a truth value.

In 1987, Ramanathan V. Guha, a mechanical engineer of Indian descent, joined the Cyc Project. While working on the Cyc Project, he developed many key parts, including the CycL representation language, and even became co-leader of the project.

At the end of 1994, Guha left the Cyc Project and soon after joined Apple Computer, Inc. While at Apple, he developed the *Meta Content Framework* (MCF), which would become the most important influence on RDF, the language of the Semantic Web. MCF is no longer used, so the most important sources of information on the framework are a set of technical reports (Guha, 1996a, 1996b, 1996c).

Let us first consider the name *Meta Content Framework*. Guha (1996c) clarifies what he means by 'meta-content':

> "We have [a] fairly clear idea of what we mean by the term content. It includes most of the documents on our hard disks, pages on the WWW, messages in email folders, etc. Meta-content is anything about this content."

In other words, 'meta-content' can be considered a synonym of the current term 'metadata', by which we mean structured data about any information object (NISO, 2004). The above quote illustrates that Guha (1996c) was primarily concerned with digital content, although further on, he does give the following example: "A book review is a piece of meta information about a piece of content – the book being reviewed."

Guha argues that there is an abundance of different types and forms of meta-content, and that a common framework is needed for representing these: "The goal of MCF is to abstract and standardize the representation of the structures we use for organizing information" (Guha, 1996b). However, he goes one step further. Given that information objects are crucially related to people, organizations and projects, MCF can describe any of these objects as well. Hence, "MCF is a general purpose structure description language." Consequently, the framework is actually broader than meta-content: it can describe not just content but *anything*. The same generality was carried through in RDF. In fact, RDF was also invented in the context of a 'metadata' initiative in the W3C, but this terminology was abandoned because its applications range much wider.

Guha presents several possible uses, most of which are related to Web usage (e.g., understanding the context and reliability of a Web page). We highlight a few examples that are remarkably reminiscent of later forms of Web usage:

- MCF could, in principle, allow for automatic creation of 'meta-pages' on the Web. If the answer to a user's query does not occur on any single page but is scattered over two or more dispersed Web pages, an MCF-based system would be able to assemble the bits and pieces into a new meta-page that contains the requested information (Guha, 1996c).

With the advent of Web 2.0, it became possible to create *mashups*, Web pages that combine elements from several sources into a new service. One key difference between Guha's idea and mashups is, however, that the latter are most often created 'by hand', whereas the (hypothetical) MCF-based mashups would be created automatically, in response to a user's query. With current technologies, this is still not possible without any human interaction, although similar hypothetical scenarios can be found in some early Semantic Web publications.

- Content directories like Yahoo! have the disadvantage that they are manually built, which is expensive and slow. Guha (1996b) advocates for jointly creating such 'maps' of Web content: "When someone in the organizations creates a page that belongs to one of the subject categories in this tree, they simply drop it under that category and every one else sees the addition. So, in a sense, this is like 'barn-raising' a Yahoo like structure."

  Barn-raising refers to a collaborative process, during which a community collectively builds a large structure, such as a barn. This concept of collaboratively building larger, labor-intensive structures has been popularized about a decade after Guha's proposal in the context of wikis, folksonomies and social bookmarking. Even the barn-raising metaphor picked by Guha is also popular with regard to wiki construction.[6]

- Guha (1996b) also proposes the idea of Personal Channels: the user specifies his/her interests in an MCF query, and "[w]henever there is a new page that matches the query, it will be delivered to the user." This is quite analogous to how RSS works, especially since nowadays it is common for services to provide RSS feeds for specific queries. This is the case for major search engines like Google (`http://www.google.com/alerts`), as well as specialized services like Web of Science or Scopus. It is interesting to note that the earliest versions of RSS are an application of RDF, which is itself a successor to MCF.

In our opinion, the main ingredient of the first two examples is *decentralization*: MCF allows for the distributed, decentralized creation of (meta)data on a large scale, which can later on be aggregated into a new structure.

MCF itself did not arise out of nowhere; there exist many clear and direct traces of CycL in MCF. Most importantly, the basic unit in MCF is the assertion. Just like in CycL, each assertion is an *n*-tuple and consists of a predicate and *n* – 1 object references. An assertion is also associated with a *layer* (essentially the same as a microtheory in Cyc) and a truth value (Guha, 1996a). The aforementioned concept of decentralization is, however, quite new: whereas Cyc constitutes a centralized knowledge base, MCF provides a framework where *anyone* can publish assertions.

At the same time that Guha worked on MCF at Apple, the World Wide Web Consortium (W3C) worked on a technical standard called PICS (Platform for Internet Content Selection; see Miller, Resnick, & Singer, 1996) that communicates ratings of a Web page's content, such that a user can choose to avoid inappropriate content (like nudity or violence). PICS was later on expanded to more generally describe Web pages, – the new version was called PICS-NG (PICS Next Generation; see Lassila, 1997). The combination of Guha's MCF and Lassila's PICS-NG, along with other metadata-related projects, culminated in a working group to consolidate these efforts. This led to the creation of the *Resource Description Framework* or RDF.

According to Miller (1998), RDF is largely a descendant of PICS-NG (MCF is mentioned only in passing). Contrary to Miller, we are convinced that some key characteristics of RDF are due to MCF, rather than PICS(-NG). We present two arguments in support of this claim. First, in spite of its name, MCF went beyond metadata and could also describe objects that were not documents and were not on the Web. Second, PICS(-NG) assumed the existence of rating services that return metadata on a given Web page. This architecture still relies on the existence of one or more centralized services. MCF's distributed nature, on the other hand, is much closer to RDF.

Semantic networks

As a network-like structure with typed links, RDF is a practical example of so-called semantic networks (Sowa, 1991). The network-like nature of RDF is emphasized by expressions such as *RDF graph* and *Web of Data*. In this section, we examine this kind of networks in more detail and pay special attention to the question where the association with semantics comes from.

Definition

A network can be defined as a structure $G = (V, E)$, consisting of a set of nodes (or vertices) V and a set of links (or edges) $E \subseteq V \times V$ that connect the nodes. What then do we mean by 'semantic networks'? There exists no universally acknowledged definition of the term 'semantic network', – indeed, it has been used in different contexts by different research fields, often referring to different entities. Henceforth, we will consider semantic networks as synonymous with *multi-relational networks* (Rodriguez & Shinavier, 2010). In this interpretation, semantic networks are an extension or generalization of single-relational networks, whose defining characteristic is that they can have more than one type of links.

Formally, semantic networks can be defined as follows. A semantic network or multi-relational network $G_m = (V, \mathbf{E})$ consists of a set of nodes V and a set of link sets $\mathbf{E} = \{E_1, \ldots, E_t\}$, where $E_j \subseteq V \times V$ $(j = 1, \ldots, t)$. Each link set in $\mathbf{E}$ represents a relation of a different type or category. There are t link sets or link types in the network.

As such, the defining characteristic of a semantic network can be likened to a nominal scale in statistics (Egghe & Rousseau, 1990), whereby observations are labeled with a name (or number). Likewise, in a semantic network the links are labeled. The label itself makes no real difference; it is just used to discern different link types. Of course, best practices dictate that clear and understandable names are to be preferred, but this is not central to semantic networks in themselves. It is, for instance, more difficult for a human being to interpret the 'e' link type in Fig. 1 than the other two types ('assailant' and 'victim'), but this makes no difference for its main function, viz. distinguishing links of this type from those of other types.

Related concepts

Semantic networks are closely related to other graph-like structures, such as concept maps, entity–relationship diagrams, topic maps etc. Sowa (2006) provides a terse overview of what he refers to as 'concept maps', although the context makes it clear that concept maps are essentially the same as semantic networks.[7] These are, he argues, a subset of the more powerful structure of 'conceptual graphs'; the latter can also incorporate quantifiers and first-order logic (see also Sowa, 2000). More generally, there exists a tension between simplicity on the one hand and expressive power on the other. Part of the attractiveness of semantic networks lies in their relative simplicity, – one does not need thorough schooling in artificial intelligence to 'get' their meaning. On the other hand, it is tempting to include useful features such as reasoning and reification. This tension persists to the present day. For instance, Notation3 or N3 (`http://www.w3.org/DesignIssues/Notation3`) is a superset of RDF that includes features such as quantifiers and entailment rules. However, it has not been adopted on a large scale. Furthermore, debate is still ongoing regarding features that should be added to or removed from RDF.

Semantic networks have been used in the fields of knowledge representation (Sowa, 1991), Artificial Intelligence, and linguistics. Apart from those, structures that are conceptually similar to semantic networks have been used in a wide variety of fields and applications. For example, many mind maps could be considered non-rigorous hierarchic semantic networks. Or, to give an example closer to the field of library and information science: thesauri as well can be considered semantic networks, with a limited number of types of relations (traditionally BT, NT, RT, USE, and UF).

The 'semantics' in semantic networks

Originally, semantic networks were network-like representations of the meaning of (typically) sentences from natural language. In a way, these early semantic networks can be considered the semantic counterparts of tree-like syntactic structures (Chomsky, 1957), which are essentially directed acyclic graphs. The primary purpose of these semantic networks was Natural Language Processing (e.g., Simmons, 1973), although their value in unambiguously recording the meaning of a given utterance – thereby facilitating discussion between researchers – was also recognized.

Fig. 1 is an example dating back to 1975, which represents the meaning of the sentence 'Dog D bit postman P'. Each node represents an object (in a broad sense); nodes are connected by binary, directed relations (arrows). Note that these relations are typed; for instance, there exists a typed relation 'victim' between situation 'B' and postman 'P'. The relation 'e' denotes the instantiation relation: for example, 'D' is an instance of the class of 'Dogs'. Thus, the semantic network makes the roles and relations of the various objects involved more explicit. Such network representations were not without controversy: “[e]ven the question of what networks have to do with semantics is one which takes some answering” (Woods, 1975).

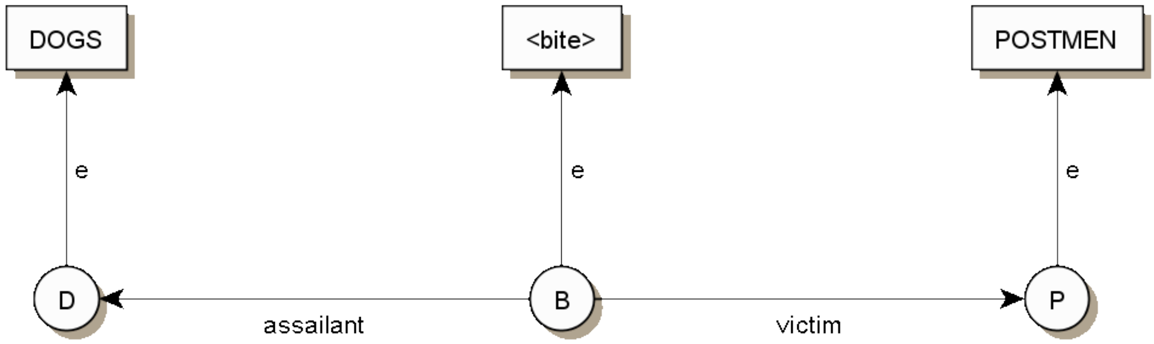

Figure 1: Example of an early semantic network, adapted from (Hendrix, 1975)

Despite early skepticism, attention shifted from rather profound theoretical issues to more practical questions regarding possible applications. This led in the 1980s and early 1990s to the development of several computer systems. The intention of many early systems was mainly to adequately represent the meaning of natural language utterances, such as the example in Fig. 1. Later systems de-emphasized the link with natural language. Instead, semantic networks became representations of concepts, relations and statements, regardless of whether they had been uttered in natural language. In a description of CASSIE, a computer system built on top of the SNePS semantic network, Shapiro and Rapaport (1986) write: “The nodes represent the objects of CASSIE’s thoughts – the things she thinks about, the properties and relations with which she characterizes them, her beliefs, etc.”

This gradual shift of focus away from language semantics may be due to the fact that it is nearly impossible for a computer to process and understand natural language without any knowledge of the world. Hence, it became a natural choice to add knowledge of the world to the data structures involved, including semantic networks. Likewise, semantic networks such as Wordnet (Fellbaum, 1998), which are concerned with lexical semantics (i.e. word meaning), are also often used as the starting point for automated reasoning systems. These examples illustrate that the direct link with semantics in a narrow sense in semantic networks was seriously diluted during the 1980s and 1990s. For many later ‘semantic networks’, the link with semantics is unclear. Nonetheless, some common characteristics did evolve. The most important ones are:

- they are schematic diagrams;
- they involve nodes and (binary) relations;
- relations are usually directed, i.e. the two participants of the relation have a different role (e.g., actor and object; action and location);
- relations are typed and thus not interchangeable.

In spite of these common elements, the interpretation of semantic networks (or, ironically, the *semantics* of semantic networks) is field- or even application-specific (Woods, 1975). Hence, more agreement is required in order to enable successful interchange between projects or applications. Rigorously defined subsets of semantic networks, such as Topic Maps (ISO, 2003) or RDF (Hayes, 2004) do provide an unambiguous interpretation and are therefore better suited to data interchange. They do, however, not solve the symbol grounding problem.

A concept can be defined as a class of similar objects that share a family resemblance; they are "actively derived from the world of objects via abstraction" (Stock, 2010, p. 1954). This definition is consistent with findings from cognitive psychology and Thomas Kuhn's philosophy (Andersen, Barker, & Chen, 1996). However, Stock (2010) also refers to individual concepts – that is, proper names. Individual concepts are a special case, in that they are instances (class members) of one or more general concepts: the individual concept *Napoleon Bonaparte* is an instance of general concepts such as *person*, *emperor*, and *father*.

Concepts can be related to other concepts. Information science distinguishes between paradigmatic and syntagmatic relations (Stock, 2010). Two concepts are syntagmatically related if they co-occur in a specific document or document part. Paradigmatic relations, on the other hand, exist independent of documents: one can, for instance, see that the concepts *apple* and *fruit* are related, irrespective of whether or not they occur together in the same document. In other words, paradigmatic relations pertain to meanings and semantic relations are "relations defined by semantic paradigms" (Murphy, 2003, p. 8). Consequently, syntagmatic relations can only be considered semantic relations in a secondary and limited sense.

An interesting example of how the use and interpretation of semantic networks may have changed can be found in Hjørland's (2007) discussion of semantic relations, defined as "the relations between concepts, meanings, or senses." He distinguishes between the concept [school] and the word 'school.' His examples are similar to the kind of relations that can be found in, for instance, Wordnet: "[School] is a kind of [educational institution]." These are paradigmatic relations between general concepts. Further on, Hjørland (2007) cites the example of a 'love' relation between two people [Tom] and [Clare]. This is remarkable for two reasons: first, [Tom] and [Clare] are individual rather than general concepts and second, the 'love' relation is *syntagmatic* instead of paradigmatic. This is exactly the kind of shift that has weakened the association between semantic networks and semantics in a linguistic sense. After all, if 'love' between [Tom] and [Clare] is a semantic relation, are there any *non*-semantic relations? Likewise, relations on the Semantic Web are both paradigmatic and syntagmatic; they can connect both general and individual concepts.

Conclusions

Many critics of the Semantic Web suggest that its name is inappropriate, since it involves semantics only at a superficial level. By tracing the origins of the (Semantic) Web, we have examined the nature of its semantics. Our overview has yielded the following insights. First, to a large extent, the association with semantics stems from the possibility of link typing. There is also a somewhat weaker association with metadata. Second, the Web has been semantically oriented virtually since its inception. Third, the Semantic Web has some obvious historical connections to Cyc and older AI research and practice.

Because of link typing, the Semantic Web can be considered an example of a semantic network. Originally network representations of the meaning of natural language utterances, semantic networks eventually refer to any networks with typed (usually directed) links. Confounding paradigmatic and syntagmatic relations yields a network that is no longer clearly semantic. In itself, it is not a new insight that the Semantic Web constitutes a network with typed links. However, to the best of our knowledge this fact has been unnoticed by previous papers that discuss the semantics of the Semantic Web.

To summarize, the Semantic Web is 'semantic' in that it relies on RDF, a graph-like data model that includes typed links. Historically, typed links have come to be associated with semantics due to their usage in semantic networks. Such networks have traditionally been used in linguistic and NLP research, and have later on become tools in different kinds of computer systems.

Acknowledgement

The present article is based on part of my dissertation (Guns, 2012). I want to thank my supervisor Ronald Rousseau for his guidance and one anonymous reviewer for some thought-provoking comments.

References

Almeida, M. B., Souza, R. R., & Fonseca, F. (2011). Semantics in the semantic web: a critical evaluation. *Knowledge Organization*, *38*(3), 187–203.

Andersen, H., Barker, P., & Chen, X. (1996). Kuhn's mature philosophy of science and cognitive psychology. *Philosophical Psychology*, *9*(3), 347–363.

Berners-Lee, T. (1989). *Information management: A proposal.* Retrieved February 22, 2012, from `http://www.w3.org/History/1989/proposal.html`

Berners-Lee, T. (2000). *Weaving the Web: The Original Design and Ultimate Destiny of the World Wide Web*. San Francisco: Harper. Paperback.

Berners-Lee, T., & Connolly, D. (1995). *RFC 1866: Hypertext Markup Language – 2.0.* IETF Request for Comments 1866. Retrieved November 29, 2012, from `http://www.w3.org/MarkUp/html-spec/html-spec_toc.html`

Berners-Lee, T., Hendler, J., & Lassila, O. (2001). The Semantic Web. *Scientific American*, *284*(5), 28–37.

Björneborn , L., & Ingwersen, P. (2001). Perspective of webometrics. *Scientometrics*, *50*(1), 65–82.

Brooks, T. A. (2002). The Semantic Web, universalist ambition and some lessons from librarianship. *Information Research*, *7*(4). Retrieved November 29, 2012, from `http://informationr.net/ir/7-4/paper136.html`

Chislenko, A. (1997). *Semantic Web vision paper.* Retrieved November 28, 2012, from `http://www.lucifer.com/~sasha/articles/SemanticWeb.html`

Chomsky, N. (1957). *Syntactic structures*. The Hague: de Gruyter Mouton.

Egghe, L., & Rousseau, R. (1990). *Introduction to Informetrics: Quantitative Methods in Library, Documentation and Information Science*. Amsterdam / New York / Oxford / Tokyo: Elsevier.

Fellbaum, C. (Ed.). (1998). *WordNet: An electronic lexical database (language, speech, and communication)* MIT Press.

Freedman, G., & Reynolds, E. G. (1980). Enriching basal reader lessons with semantic webbing. *Reading Teacher, 33*(6), 677–84.

Gärdenfors, P. (2004). How to make the semantic web more semantic. In A. C. Varzi & L. Vieu (Eds.), *Formal Ontology in Information Systems: Proceedings of the Third International Conference* (Vol. 114, pp. 17–34). IOS Press.

Guha, R. V. (1996a). *Meta Content Framework* (Apple Technical Report #167). Apple Computer. Retrieved November 28, 2012, from http://www.guha.com/mcf/mcf_spec.html

Guha, R. V. (1996b). *Meta Content Framework: A whitepaper* (Apple Technical Report #168). Apple Computer. Retrieved November 28, 2012, from http://www.guha.com/mcf/wp.html

Guha, R. V. (1996c). *Towards a theory of meta-content* (Apple Technical Report #169). Apple Computer. Retrieved November 28, 2012, from http://www.guha.com/mcf/mc.html

Guns, R. (2012). *Missing links: Predicting interactions based on a multi-relational network structure with applications in informetrics*. Unpublished doctoral dissertation, University of Antwerp.

Harnad, S. (2002). *The Syntactic Web.* Retrieved May 20, 2012 from http://lists.w3.org/Archives/Public/www-archive/2002Sep/0114.html

Hayes, P. (2004). *RDF semantics.* W3C Recommendation. Retrieved November 28, 2012, from http://www.w3.org/TR/rdf-mt/

Heflin, J. (2001). *Towards the semantic web: Knowledge representation in a dynamic, distributed environment*. Unpublished doctoral dissertation, University of Maryland.

Heflin, J., & Hendler, J. (2000). Semantic interoperability on the web. In *Proceedings of Extreme Markup Languages 2000* (pp. 111–120). Alexandria, VA: Graphic Communications Association. Retrieved November 28, 2012 from http://www.cse.lehigh.edu/~heflin/pubs/extreme2000.pdf

Hendrix, G. G. (1975). *Expanding the utility of semantic networks through partitioning* (Tech. Rep.).

Hickson, I. (Ed.). (2012a). *HTML5: A vocabulary and associated APIs for HTML and XHTML*. W3C Working Draft 29 March 2012. Retrieved November 28, 2012, from http://www.w3.org/TR/2012/WD-html5-20120329/

Hickson, I. (Ed.) (2012b). *HTML microdata*. W3C Working Draft 25 October 2012. Retrieved November 28, 2012, from http://www.w3.org/TR/2012/WD-microdata-20121025/

Hjørland, B. (2007). Semantics and knowledge organization. *Annual Review of Information Science & Technology*, *41*(1), 367–405.

Horrocks, I., Patel-Schneider, P. F., & van Harmelen, F. (2003). From SHIQ and RDF to OWL: the making of a web ontology language. *Web Semantics: Science, Services and Agents on the World Wide Web*, *1*(1), 7–26.

ISO. (2003). *ISO/IEC 13250: Topic Maps* (2nd edition).

Khare, R., & Çelik, T. (2006). Microformats: A pragmatic path to the Semantic Web. In *Proceedings of the 15th International Conference on World Wide Web* (pp. 865–866). New York: ACM.

Kogan, Y., Michaeli, D., Sagiv, Y., & Shmueli, O. (1998). Utilizing the multiple facets of WWW contents. *Data & Knowledge Engineering*, *28*(3), 255–275.

Lassila, O. (1997). *PICS-NG Metadata Model and Label Syntax* (Tech. Rep.). W3C Recommendation. Retrieved November 28, 2012 from http://www.w3.org/TR/NOTE-pics-ng-metadata

Legg, C. (2007). Ontologies on the Semantic Web. *Annual Review of Information Science & Technology*, *41*(1), 407–451.

Lenat, D. B. (1995). CYC: a large-scale investment in knowledge infrastructure. *Communications of the ACM*, *38*(11), 33–38.

Lenat, D. B., & Guha, R. V. (1991). The evolution of CycL, the Cyc representation language. *ACM SIGART Bulletin*, *2*(3), 84–87.

Luke, S., Spector, L., & Rager, D. (1996). Ontology-based knowledge discovery on the World-Wide Web. In A. Franz & H. Kitano (Eds.), *Working Notes of the Workshop on Internet-Based Information Systems at the 13th National Conference on Artificial Intelligence (AAAI96)* (pp. 96–102). Menlo Park: AAAI Press.

Miller, E. (1998). An introduction to the Resource Description Framework. *D-Lib Magazine*, *4*(5). Retrieved November 30, 2012, from http://www.dlib.org/dlib/may98/miller/05miller.html

Miller, J., Resnick, P., & Singer, D. (1996). *Rating Services and Rating Systems (and Their Machine Readable Descriptions)* (Tech. Rep.). W3C Recommendation. Retrieved November 28, 2012 from http://www.w3.org/TR/REC-PICS-services/

Murphy, M. L. (2003). *Semantic Relations and the Lexicon: Antonymy, Synonymy and other Paradigms*. Cambridge: University Press.

NISO. (2004). *Understanding metadata*. Bethesda, MD. Retrieved November 28, 2012 from http://www.niso.org/publications/press/UnderstandingMetadata.pdf

O'Hara, K., & Hall, W. (2010). Semantic Web. In M. J. Bates & M. N. Maack (Eds.), *Encyclopedia of library and information science* (Third ed., pp. 4663–4676). CRC Press.

Rodriguez, M. A., & Shinavier, J. (2010). Exposing multi-relational networks to single-relational network analysis algorithms. *Journal of Informetrics*, *4*(1), 29–41.

Rousseau, B., & Rousseau, R. (2002). Some ideas concerning the semantic web. *Library and Information Service*, *8*, 39–49.

Shapiro, S. C., & Rapaport, W. J. (1986). SNePS considered as a fully intensional propositional semantic network. In *AAAI-86 Proceedings* (pp. 278–283).

Shirky, C. (2003). *The Semantic Web, syllogism, and worldview.* Retrieved January 21, 2012, from http://www.shirky.com/writings/semantic_syllogism.html

Simmons, R. F. (1973). Semantic networks: Their computation and use for understanding english sentences. In R. C. Schank & K. M. Colby (Eds.), *Computer Models of Thought and Language.* San Francisco: Freeman.

Sowa, J. F. (Ed.). (1991). *Principles of Semantic Networks: Explorations in the Representation of Knowledge*. San Mateo, CA: Morgan Kaufmann.

Sowa, J. F. (2000). *Knowledge Representation: Logical, Philosophical, and Computational foundations*. Brooks/Cole. Hardcover.

Sowa, J. F. (2006). Concept mapping. In *Proceedings of the AERA Annual Meeting, 10 April 2006.* San Francisco. Retrieved November 30, 2012, from `http://www.jfsowa.com/talks/cmapping.pdf`

Sparck-Jones, K. (2004). What's new about the Semantic Web? Some questions. *ACM SIGIR Forum*, 38, 18–23.

Stock, W. G. (2010). Concepts and semantic relations in information science. *Journal of the American Society for Information Science and Technology*, *61*(10), 1951–1969.

Thelwall, M. (2004). *Link analysis: An information science approach* (1st ed.). Emerald Group Publishing.

Thelwall, M., Vaughan, L., & Björneborn , L. (2006). Webometrics. *Annual Review of Information Science and Technology*, *39*(1), 81–135.

Uschold, M. (2003). Where are the semantics in the Semantic Web? *AI Magazine*, *24*(2), 25–36.

Woods, W. A. (1975). What's in a link: Foundations for semantic networks. In D. G. Bobrow & A. Collins (Eds.), *Representation and understanding: Studies in cognitive science.* New York / San Francisco / London: Academic Press.

---

[1] http://www.w3.org/DesignIssues/Topology.html#4

[2] http://www.w3.org/Talks/INET95/Overview.html

[3] http://code.google.com/webstats/2005-12/linkrels.html

[4] http://www.opencyc.org

[5] http://cyc.com/cyc/technology/whatiscyc_dir/whatsincyc

[6] See, e.g., http://www.wikipatterns.com/display/wikipatterns/BarnRaising or http://meatballwiki.org/wiki/BarnRaising.

[7] It is unclear where exactly Sowa draws the line. On p. 7 he gives the following examples of 'relational graphs': "Concept Maps, Topic Maps, RDF, OWL, Type Hierarchies, Entity-Relationship Diagrams, frame systems, and many semantic networks." We prefer to consider semantic networks as the more general concept.