

The Semantic Web: The Origins of Artificial Intelligence Redux

Harry Halpin

ICCS, School of Informatics

University of Edinburgh

2 Buccleuch Place

Edinburgh EH8 9LW

Scotland UK

Fax:+44 (0) 131 650 458

E-mail:h.halpin@ed.ac.uk

Corresponding author is Harry Halpin.

For further information please contact him.

This is the tear-off page.

To facilitate blind review.

1 Introduction

The World Wide Web is considered by many to be the most significant computational phenomenon yet, although even by the standards of computer science its development has been chaotic. While the promise of artificial intelligence to give us machines capable of genuine human-level intelligence seems nearly as distant as it was during the heyday of the field, the ubiquity of the World Wide Web is unquestionable. If anything it is the Web, not artificial intelligence as traditionally conceived, that has caused profound changes in everyday life. Yet the use of search engines to find knowledge about the world is surely in the spirit of Cyc and other artificial intelligence programs that sought to bring all world knowledge together into a single database. There are, upon closer inspection, both implicit and explicit parallels between the development of the Web and artificial intelligence.

The Semantic Web effort is in effect a revival of many of the claims that were given at the origins of artificial intelligence. In the oft-quoted words of George Santayana, “those who do not remember the past are condemned to repeat it.” There are similarities both in the goals and histories of artificial intelligence and current developments of the Web, and in their differences the Web may find a way to escape repeating the past.

2 The Hilbert Program for the Web

The development of World Wide Web has become a field of endeavor in itself, with important ramifications for the world at large, although these are noticed more by industry than philosophy. The World Wide Web is thought of as a purely constructed system; its problems can be construed as engineering problems rather than as scientific, mathematical, or philosophical problems. The World Wide Web problems grew dramatically with its adoption, and during the “browser wars” between Netscape and Microsoft, it was feared that the Web would fragment as various corporations created their own proprietary extensions to the Web. This would defeat the development of the original purpose of the Web as a universal information space. In response to this crisis, Tim Berners-Lee, the inventor of the Web, formed a non-profit consortium called the World Wide Web Consortium (W3C) that “by promoting interoperability and encouraging an open forum for discussion” will lead “the technical evolution of the Web” by its three design principles of interoperability, evolution, and decentralization (W3C, 1999). Tim Berners-Lee is cited as the inventor of the Web for his original proposal for the creation of the Web in 1989, his implementation of the first web browser and server, and his initial specifications of URIs, HTTP, and HTML (2000). Due to this, the W3C was joined by a wide range of companies, non-profits, and academic institutions (including Netscape and Microsoft), and through its

working process managed to both halt the fragmentation of the Web and create accepted Web standards through its consensus process and its own research team. The W3C set three long-term goals for itself: universal access, Semantic Web, and a web of trust, and since its creation these three goals have driven a large portion of development of the Web (W3C, 1999).

One comparable program is the Hilbert Program in mathematics, which set out to prove all of mathematics follows from a finite system of axioms and that such an axiom system is consistent (Hilbert, 1922). It was through both force of personality and merit as a mathematician that Hilbert was able to set the research program and his challenge led many of the greatest mathematical minds to work. The Hilbert Program shaped irrevocably the development of mathematical logic for decades, although in the end it was shown to be an impossible task. In a similar fashion, even if the program of Berners-Lee and the W3C fails (although by its more informal nature it is unlikely to fail by a result as elegant as the Second Incompleteness Theorem), it will likely produce many insights into how the Web may, in the words of Berners-Lee, “reach its full potential” (2000).

At first, the W3C was greeted with success, not only for standardizing HTML, but also for the creation of XML, an extensible markup language that generalized HTML so that anyone could create their own markup language as long as they followed a syntax of tree-structured documents with links. While originally created to separate presentation from content, it soon became used primarily to move data of any sort across the Web, since “tree-structured documents are a pretty good transfer syntax for just about anything,” combined with the weight given to XML by the W3C’s official recommendation of it as a universal standard (Thompson, 2001). XML is poised to become a universal syntax, an “ASCII for the 21st century.” Immediately following the prospects for “moving beyond syntax to semantics” arose (Thompson, 2001). This is where the next step in the W3C vision appears: the Semantic Web, defined by Berners-Lee as “an extension of the current Web in which information is given well-defined meaning, enabling computers and people to work in better cooperation” (2001). Berners-Lee continued that “most of the Web’s content today is designed for humans to read, not for computer programs to manipulate meaningfully” and so the Semantic Web must “bring structure to the meaningful content of Web pages, creating an environment where software agents roaming from page to page can readily carry out sophisticated tasks for users” (2001). This vision, implemented in knowledge representation, logic, and ontologies is strikingly similar to the vision of artificial intelligence.

3 Brief History

3.1 Artificial Intelligence

To review the claims of artificial intelligence in order to clarify their relation to the Semantic Web, we are best served by remembering the goal of AI as stated by John McCarthy at

the 1956 Dartmouth Conference: "The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it"(McCarthy et al., 1955). However, "intelligence" itself is not clearly defined. The proposal put forward by McCarthy gave a central role to "common-sense," so that "a program has common sense if it automatically deduces for itself a sufficient wide class of immediate consequences of anything it is told and what it already knows"(McCarthy, 1959). A plethora of representation schemes, ranging from semantic networks to frames, all flourished to such an extent that Herbert Simon wrote that "machines will be capable, within twenty years, of doing any work that a man can do"(1965). While many of these programs, from Logic Theorist(Simon and Newell, 1958) to SHRDLU(Winograd, 1972) managed to simulate intelligence in a specific domain such as proving logical theorems or moving blocks, it became clear that this strategy was not scaling up to the level of general intelligence. Although AI had done well in "tightly-constrained domains," extending this ability had "not proved straightforward"(Winston, 1976). Even within a specific knowledge representation form such as semantic networks, it was shown that a principal element such as a link was interpreted in at least three different ways(Woods, 1975). Knowledge representations were not obviously denoting the knowledge they supposedly represented. This led to a general reaction to give a formal account of the knowledge using a well-understood framework, such as first-order predicate logic, which was equivalent to most of the knowledge representation systems used at the time(Hayes, 1977). The clear next step was the formalization of as much common-sense knowledge as possible using rigorous standards of logic, in order to overcome small, domain-specific strategies(Hayes, 1986). Yet, these approach seemed to never converge on a universal formal manner of representing all knowledge, as revealed by the influential Brachman-Smith survey(Brachman and Smith, 1991). This survey was a testament to an immense range and diversity of AI systems, a virtual Tower of Babel. Unification and formalization of all "common-sense" knowledge seemed even further away. While some remaining AI researchers maintained that all of necessary common-sense knowledge could be encoded shortly(Lenat and Feigenbaum, 1987), many other researchers left the field and the AI industry collapsed. To this day Lenat is still encoding "common-sense" into Cyc(Lenat, 1990). Brian Smith published an oft-overlooked critique of the entire research program to formalize common-sense, noting that all useful knowledge was situated in the particular task at hand and the agent, and it seemed unlikely that any traditional knowledge representation or logical foundation could capture these aspects of knowledge(1991). If this was true, the claim of upholding that a machine could simulate human level intelligence through sheer formalization of facts and inferences seemed doomed, although such a program might produce useful technology regardless of the original claims of AI. Instead of formalizing common-sense, Smith instead asked what lessons artificial intelligence could learn from indexing and retrieving information in a worldwide digital collection of documents, a statement that strangely prefigures the

development of the Web.

3.2 The Semantic Web

The Web is returning to the traditional grounds of artificial intelligence in order to solve its own problems. It is a mystery to many why Berners-Lee and others believe the Web needs to transform into the Semantic Web. However, it may be necessitated by the growing problems of information retrieval and organization. The first incarnation of the Semantic Web was meant to address this problem by encouraging the creators of web-pages to provide some form of metadata (data about data) to their web-page, so simple facts like identity of the author of a web-page could be made accessible to machines. This approach hopes that people, instead of hiding the useful content of their web pages within text and pictures that was only easily readable by humans, would create machine-readable metadata to allow machines to access their information. To make assertions and inferences from metadata, inference engines would be used. The formal framework for this metadata, called the Resource Description Framework (RDF), was drafted by Hayes, one of the pioneers of artificial intelligence. RDF is a simple language for creating assertions about propositions(Hayes, 2004). The basic concept of RDF is that of the "triple": any statement can be composed into a subject, a predicate, and the object. "The creator of the web-page is Henry Thompson" can be phrased as `www.inf.ed.ac.uk/ht dc:creator "Henry Thompson"`. The framework was extended to that of a full ontology language as described by a description logic. This Web Ontology Language (OWL) is thus more expressive than RDF(Welty et al., 2004). The Semantic Web paradigm made one small but fundamental change to the architecture of the Web: a resource (that is, anything that can be identified by a URI) can be about anything. This means that URIs, that were formerly used to denote mostly web-pages and other data that has some form of byte-code on the Web can now be about anything from things whose physical existence is outside the Web to abstract concepts(Jacobs and Walsh, 2004). A URI can denote not just a web-page about the Eiffel Tower but the Eiffel Tower itself (including if there is no web-page at that location) or even a web-page about "the concept of loyalty." This change is being reworked by Berners-Lee into the revised URI specification and an upcoming normative W3C document entitled "The Architecture of the Web"(Jacobs and Walsh, 2004). What was at first manually annotating web-pages with proposition-like metadata now can become the full-scale problem of knowledge representation and ontology development, albeit with goals and tools that have been considerably modified since their inception during the origins of artificial intelligence. The question is has the Semantic Web learned anything from artificial intelligence?

4 Differences of the Semantic Web

The first major difference between early artificial intelligence and the Semantic Web is that the Semantic Web is clearly not pursuing the original goal of AI as stated by the Dartmouth

Proposal: “human-level intelligence”(McCarthy et al., 1955). The goal of the Semantic Web is more modest and in line with later artificial intelligence research, that of creating machines capable of exhibiting “intelligent” behavior. This goal is much harder to test, since if “intelligence” for machines is different than humans intelligence, there exists no similar Turing Test to detect merely machine-level intelligence (Turing, 1950). However, there are reasons that the Semantic Web engineers have for their hope that their project might fulfill some of the goals of artificial intelligence, in particular the goal of creating usable ontologies of the real world.

4.1 Difference of scale

For the first time in human history, truly mammoth amounts of raw information is available for transformations into ontologies. While it is unclear exactly how much data is on the Web, there is more human-readable data in digital form than ever before, and increasing demand for more intelligent ways of navigating and organizing it. This contrasts with the origins of artificial intelligence, where much of the knowledge base that were in digital form was quite small. Although previous work on the inability of domain-specific AI to scale hint that sheerly increasing the amount of information may not help(Winston, 1976), increasing scale might help. Even if it does, the information does not scale to a general database of common-sense relations and the like, the domain-specific knowledge available in digital form is also much larger than those available previously. Since most of this information is available only in human-readable form or in traditional databases, a rapidly growing body of work attempts to address the automatic extraction of metadata and ontologies from web-pages (Dill et al., 2003). One Semantic Web project, Friend-of-a-Friend (FOAF), boasts over a million users.¹ The sheer quantity of human-made ontologies and metadata available over the web, while still definitely in its infancy and taking longer to become popularly adopted than envisioned, definitely gives the potential for the economies of scale of the Semantic Web to be larger than that of artificial intelligence.

4.2 Description Logic

As noted earlier, one problem with traditional artificial intelligence was the lack of an agreed upon formal foundation with well-described and understood properties. Usually ontologies were created by small research groups, with each group having its own form of knowledge representation, although almost all representational schemes were found to be equivalent to first-order logic. Since that time, a part of the artificial intelligence community has developed description logic, which is a subset of first-order logic with well-known properties. Unlike first-order logic, description logics have proved to be decidable and of a tractable complexity class(Borgida, 1996). These logics are well-studied in both academic and industrial use as well, with OWL closely modeling itself on

¹<http://www.foaf-project.org/>

the CLASSIC project(Borgida et al., 1989) and its descendants. The W3C has agreed upon the use of description logics and “triples” as its guiding principles for web ontologies. However, description logics highly constrain what one can say in order to maintain decidable inference and this can lead to a language that may be too restrictive to say many ordinary logical statements that could be made about the Web(Hayes, 2002). OWL Full goes beyond some of these limits of description logic, but since any flavor of OWL has yet to see widespread use, it is difficult to say how desirable decidability can be. For RDF, the W3C has chosen to stay to a simple propositional calculus, with every statement encoded as a simple assertion(Hayes, 2004).

4.3 Decentralized

Ontologies have been developed by a single core group of people, whether an academic research group or a particular company. This mode of development is centralized by nature. The Semantic Web allows decentralized creation of ontologies, hoping that industries and researchers will reach consensus on large-scale ontologies. In the spirit of MYCIN(Shortliffe, 1976), the life sciences have been one of the first domains to begin standardizing ontologies such as GeneOntology,² and this ontology can coexist and be used with similar ones such as BioPax.³ Ontologies can also be explicitly mapped to each other. These ontologies might remain mutually incommensurable except for human-created bridges. So, automated creation of these mappings is still an active and difficult area of research(Bouquet et al., 2003). Still, there is sign of success even a heavily decentralized metadata creation, such as the Friend Of A Friend metadata project. It uses the hand-coded work of many decentralized groups of people to create a truly huge, if simple, network of metadata that map people and their interests.

4.4 Universality

Although many traditional knowledge representation systems claimed to be universal, the ability for any component of a Semantic Web ontology to be given an universally unique name gives the Web a distinct advantage. “The Semantic Web, in naming every concept simply by a URI, lets anyone express new concepts that they invent with minimal effort”(Berners-Lee et al., 2001). This mechanism gives ontologies the means to be accessed from anywhere with web access, and made accessible by simply putting their ontology on the Web.

4.5 Open-World Assumption

One further note upon the development of the Semantic Web, although this is unclear whether this a distinct advantage, is that while traditional AI systems operated by a “closed-world” assumption, the Semantic Web operates by an “open-world” assumption, restricting itself to monotonic reasoning(Hayes, 2001). The reason for this, is that on the Web

²<http://www.geneontology.org/>

³<http://www.biopax.org/>

reasoning “needs to always take place in a potentially open-ended situation: there is always the possibility that new information might arise from some other source, so one is never justified in assuming that one has ‘all’ the facts about some topic”(Hayes, 2001).

5 Unsolved Problems in the Artificial Intelligence

As much as the Semantic Web effort has made careful and web-scale improvements over the foundations of knowledge representations used traditionally in artificial intelligence, it also inherits some of the more dangerous problems of artificial intelligence. These must be at least recognized for the Semantic Web, otherwise it returns the problems of AI “the first time as tragedy, the second as farce”(Marx, 1852).

5.1 The Knowledge Representation Problem

In particular, it inherits what I term *Knowledge Representation Problem*. If knowledge representations are fundamentally stand-in surrogates for facets of the world, then “how close is the surrogate to the real thing? What attributes of the original does it capture and make explicit, and which does it omit? Perfect fidelity is in general impossible, both in practice and in principle. It is impossible in principle because any thing other than the thing itself is necessarily different from the thing itself.” This leads to the conclusion that “imperfect surrogates mean incorrect inferences are inevitable”(Davis et al., 1993). The scale of the Semantic Web may aggravate instead of solve the problem. With a decentralized method of creating knowledge representations, it becomes increasingly difficult to guess what features of the world people might formalize into an ontology. This will lead to many ontologies that are about the same things, yet it will be unable to tell if the elements of two ontologies are equivalent. Even if there was unambiguous human-understandable documentation that showed two ontology elements to be equivalent, the task of mapping between many small ontologies manually is immense. One way to resolve this would be to use only a few well-specified large ontologies, yet one loses the ability to map one’s locally rich semantic space to a custom ontology. It is also hard to tell how “brittle” these ontologies are. This is reminiscent of the problem of domain-specific AI systems being unable to scale. To be automated, overcoming this problem requires at least non-monotonic reasoning or at most the original goal of AI, human-level intelligence(McCarthy et al., 1955).

5.2 The Higher-Order Problem

This problem occurs when a logical system tries to make inferences about its own contents. This problem leads to predicates *about predicates*, with the possibility of quantification over already quantified predicates. This transforms predicate logic into higher-order logic, which has less well-known and definitely less tractable properties. In an attempt to solve the

problem of attribution, it is often considered useful to employ reification of RDF statements. However, this has been found to both computationally difficult to implement and lead to misleading attributions. As stated by the RDF Semantics, “Since an assertion of a reification of a triple does not implicitly assert the triple itself, this means that there are no entailment relationships which hold between a triple and a reification of it’, and so making it very difficult to fit reified statements into the model theory”(Hayes, 2004). This problem had been discovered in AI by the work in computational reflection and reification(Smith, 1984).

5.3 The Abstraction Problem

Abstraction is both a benefit and a curse for the Semantic Web, especially when classes and individuals are introduced by OWL. The question about whether to implement a knowledge representation as either abstract or concrete is subtle(Smith, 1996). For example, the “Dartmouth School of Art” can be thought of as an concrete instance of the class of all schools, or as a abstract class which remains the same regardless of the moving of the physical building or the change of staff. It then becomes unclear what one is referring in statements such as “The Dartmouth School of Art is now specializing in sculpture” or “The Dartmouth School of Art has changed its address.” This problem is recognized by the OWL ontology group. The OWL documentation mentions both that “in certain contexts something that is obviously a class can itself be considered an instance of something else” and “it is very easy to confuse the instance-of relationship with the subclass relationship”(Welty et al., 2004). This makes ontology mapping and merging exceedingly difficult. While the ability to divide the world into classes and instances provide description logics with a set of principles, it does not make mapping between what one person considers a class and another considers an instance straightforward.

5.4 The Frame Problem

The question of how to represent time in an open world is another question from artificial intelligence that haunts the Semantic Web. RDF attempts to avoid this problem by stating that “does not provide any analysis of time-varying data”(Hayes, 2004). Yet, it would seem that any statement about an URI is not meant to last forever, especially as URIs and their contents have a tendency to change. Berners-Lee attempts to avoid this problem in a note “Cool URIs don’t Change,”⁴ in which he notes that the change of a URI damages its ability to be universally linked and have statements made about it. However, despite this principle being made fundamental in new Web standards, it at the current moment does not stand true about the Web and we have no reason to believe that it will soon in the future(Jacobs and Walsh, 2004). There is already a need to make temporally-qualified statements using metadata and ontologies. However, as pointed out by the Frame Problem, the issue of handling assumptions

⁴<http://www.w3.org/Provider/Style/URI>

about time in artificial intelligence has proven remarkably difficult to formalize (McCarthy and Hayes, 1969). Their example is that if “we had a number of actions to be performed in sequence we would have quite a number of conditions to write down that certain actions do not change the values of certain fluents”(McCarthy and Hayes, 1969). There is no agreed upon model of time with properties that are well understood. In fact, there are many theories of time with contradictory properties(Hayes, 1995).

5.5 The Symbol Grounding Problem

This problem is stated as “How can the semantic interpretation of a formal symbol system be made intrinsic to the system, rather than just parasitic on the meanings in our heads? How can the meanings of the meaningless symbol tokens, manipulated solely on the basis of their (arbitrary) shapes, be grounded in anything but other meaningless symbols?”(Harnad, 1990). One answer is to map it via formal semantics to a model theory. However, although the model may resemble the part of the world that it models, it could also model it in a limited fashion due to the Knowledge Representation problem. Therefore, it is needed to “ground” the symbols in some real world object(Davis et al., 1993). It is difficult to imagine how what this would practically entail, perhaps some form of sensors with direct causal content as Harnad suggests(1990). However, it is not clear that such a direct connection is needed, for even humans do not remain in constant causal contact with their subject matter. In fact, this ability to connect and disconnect our representations with their subject matter is a reason for the origins of intentionality and representations in humans(Smith, 1996). A machine must use whatever information it can find out about the subject matter, even if that information is by nature partial. Although direct causal contact is limited by machines on the Web to those things that are Web accessible, it could include various statements (and entailments) encoded in the Semantic Web by humans and the immense amount of content created on the Web by human users about real world content. This gives the Semantic Web the ability to skirt around the problem by using Web accessible information that is grounded in human authority and human sensory contact with the world outside the Web, although it is far from a satisfactory solution.

5.6 The Problem of Trust

This problem is virtually non-existent in AI since most knowledge representation systems were created by small groups who trusted their members. With the decentralization of ontology creation and the ability for ontologies to universally import, use, and perhaps map and merge to each other, there is a real need to know if the creator of some ontology is trustworthy. This leads to serious issues with ontology evaluation and has been one of the reasons concepts like reification in RDF were originally pursued.

5.7 Engineering or Epistemology?

The Semantic Web may not be able to solve many of these problems. Many Semantic Web researchers pride themselves on being engineers as opposed to artificial intelligence researchers, logicians, or philosophers, and have been known to believe that many of these problems are engineering problems. While there may be suitable formalizations for time and ways of dealing with higher-order logic, the problems of knowledge representation and abstraction appear to be epistemological characteristics of the world that are ultimately resistant to any solution. It may be impossible to solve some of these problems satisfactorily, yet having awareness of these problems can only help the development of the Web.

6 Conclusions

6.1 The Web as Universal Computing

The Web could do something even more interesting than that what the Semantic Web promises. Many in industry are interested in “Web Services” currently, which consists in using the Web to send data across the Web.⁵ This rather mundane observation, if taken to its logical conclusion, has serious ramifications. Due to its characteristics of universality, it could implement what I call “Turing’s Promise.” Turing created a universal abstract model for computers in the form of his universal Turing machines(1936). While all computers are realizations of universal Turing machines, actual computers in practice are a wide variety of incompatible hardware and software, and so not universal in actual use. XML provides a universal syntax for data, and URIs provide a universal way of naming things. A theoretical programming language that takes URIs as its base naming convention and uses some version of XML and RDF as its core data structures and typing systems could qualify as a universal programming language, one that by nature is no longer constrained by the von Neumann style(Backus, 1978). The creation of an universal way of handling data on the Web through a functional language is an avenue that has not yet been explored. Implementing this language would lead to the Web being not one large knowledge representation system, but a distributed universal computer that can take advantage of the universal information space that is the Web.

6.2 Redeeming Turing’s Promise

As regards the prospects for artificial intelligence, the Web deserves the attention of both practitioners and historians as it exhibits a wide variety of features that build upon long-standing problems. The Web presents a widely-used and differing architecture than traditional computers, and so new initiatives in logic and semantics are needed. In the latest Semantic Web initiative, the W3C is building upon lessons learned at the origins of artificial intelligence, and yet if anything the problems posed and discovered by artificial intelligence are more pressing now than ever. The original promise

⁵<http://www.w3.org/2002/ws/>

of artificial intelligence have been lessened in ambition yet made more trenchant due to the need to operate over a universal information space. This should not be surprising: neither intelligence nor universality are trivial, and only a detailed examination of the past and a sharp eye in the present will help the Web succeed, redeeming Turing's original promise, if not of artificial intelligence, of universal computation.

References

- Backus, J. (1978). Can programming be liberated from the von Neumann style?: a functional style and its algebra of programs. *Communications of the ACM*, 21(8).
- Berners-Lee, T. (2000). *Weaving the Web*. Texere Publishing, London.
- Berners-Lee, T., Hendler, J., and Lassila, O. (2001). The Semantic Web. *Scientific American*.
- Borgida, A. (1996). On the relative expressiveness of description logics and predicate logics. *Artificial Intelligence*, 82.
- Borgida, A., Brachman, R., McGuinness, D., and Resnick, L. (1989). CLASSIC: A structural data model for objects. In *Proceedings of the 1989 ACM SIGMOD International Conference on Management of Data*.
- Bouquet, P., Serafini, L., and Zanobini, S. (2003). Semantic coordination: A new approach and an application. In *International Semantic Web Conference*.
- Brachman, R. and Smith, B. (1991). Special issue on knowledge representation. *SIGART Newsletter*, 70.
- Davis, R., Shrobe, H., and Szolovits, P. (1993). On the relative expressiveness of description logics and predicate logics. *AI Magazine*, 14(1).
- Dill, S., Eiron, N., Gibson, D., and et al. (2003). Semtag and Seeker: Bootstrapping the Semantic Web via automated semantic annotation. In *Proceedings of the International World Wide Web Conference*.
- Harnad, S. (1990). The Symbol Grounding Problem. *Physica*.
- Hayes, P. (1977). In defense of logic. In *Proceedings of International Joint Conference on Artificial Intelligence*, pages 559–565.
- Hayes, P. (1986). The second naive physics manifesto. In *Formal Theories of the Commonsense World*. Ablex.
- Hayes, P. (1995). A catalog of temporal theories. Technical report, University of Illinois. Tech report UIUC-BI-AI-96-01.
- Hayes, P. (2001). Why must the web be monotonic? Technical report, IHMC. <http://lists.w3.org/Archives/Public/www-rdf-logic/2001Jul/0067.html>.
- Hayes, P. (2002). Catching the dream. Technical report, IHMC. <http://www.aifb.uni-karlsruhe.de/sst/is/WebOntologyLanguage/hayes.htm/>.
- Hayes, P. (2004). RDF Semantics. Technical report, W3C. <http://www.w3.org/TR/2004/REC-rdf-mt-20040210/>.
- Hilbert, D. (1922). *Neubegründung der Mathematik: Erste Mitteilung, Abhandlungen aus dem Seminar der Hamburgischen Universität*. Mathematisches Institut, Universität Göttingen.
- Jacobs, I. and Walsh, N. (2004). Architecture of the World Wide Web. Technical report, W3C. <http://www.w3.org/TR/webarch/>.
- Lenat, D. (1990). Cyc: Towards Programs with Common Sense. *Communications of the ACM*, 33(8):30–49.
- Lenat, D. and Feigenbaum, E. (1987). On the Thresholds of Knowledge. In *Proceedings of International Joint Conference on Artificial Intelligence*.
- Marx, K. (1852). *The Eighteenth Brumaire of Louis Bonaparte*.
- McCarthy, J. (1959). Programs with common-sense. *Nature*, 188:77–91.
- McCarthy, J. and Hayes, P. (1969). Some philosophical problems from the standpoint of Artificial Intelligence. In *Machine Intelligence*, volume 4.
- McCarthy, J., Minsky, M., Rochester, N., and Shannon, C. (1955). A Proposal for the Dartmouth Summer Research Project on Artificial intelligence. Technical report.
- Shortliffe, E. (1976). *MYCIN: Computer-based Medical Consultations*.
- Simon, H. (1965). *The Shape of Automation for Men and Management*.
- Simon, H. and Newell, A. (1958). Heuristic problem solving: The next advance in operations research. *Operations Research*, 6.
- Smith, B. C. (1984). Reflection and semantics in LISP. *Proceedings of 11th ACM SIGACT-SIGPLAN symposium on Principles of programming languages*, pages 23–35.
- Smith, B. C. (1991). The Owl and the Electric Encyclopedia. *Artificial Intelligence*, 47:251–288.
- Smith, B. C. (1996). *On the Origin of Objects*. MIT Press, Cambridge, Massachusetts.
- Thompson, H. S. (2001). Putting XML to Work. Technical report, University of Edinburgh.
- Turing, A. M. (1936). On Computable Numbers, with an Application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society*, 42.
- Turing, A. M. (433-460). Computing machinery and intelligence. *Mind*, 59.
- W3C (1999). W3C Mission Statement. Technical report. <http://www.w3.org/Consortium/>.
- Welty, C., Smith, M., and McGuinness, D. (2004). OWL Web Ontology Language Guide. Technical report, W3C. <http://www.w3.org/TR/2004/REC-owl-guide-20040210>.
- Winograd, T. (1972). Procedures as a Representation for Data in a Computer Program for Understanding Natural Language. *Cognitive Psychology*, 3(1).
- Winston, P. (1976). AI memo no. 366. Technical report, MIT.

Woods, W. (1975). What's in a link: Foundations for semantic networks. In *Representation and Understanding: Studies in Cognitive Science*, pages 35–82.