From Open IS Semantics to the Semantic Web: The Road Ahead

David G. Schwartz, *Bar-Ilan University*

The Semantic Web should ideally combine the best of AI research with the best of Internet research. This analysis looks at distributed AI, Open Information Systems Semantics, and the Semantic Web and presents several open research issues.

Sharing knowledge between diverse heterogeneous sources has been at the top of the AI community agenda for over 40 years, since the early days of blackboard systems research and semantic-information processing. The appeal and allure of automated understanding and processing of rich contextual knowledge rather than dry data has surfaced in different guises in each of the last four decades (see the "History of Semantic Reasoning Systems" sidebar). Many of these initiatives sought to encode knowledge in some machine-readable and communicable form for use across system boundaries for the eventual use of multiple agents or participants.

In between all that AI activity, the Internet emerged and took center stage as a unifying infrastructure promising unprecedented connectivity, if not unprecedented semantic agreement.

And now there's the Semantic Web.

In the years since its introduction as an Internet application for sharing data,[1] the Web has grown explosively, metamorphosing toward storing, transmitting, and processing semantic data. The Semantic Web has great potential significance in the continuum of AI research. My goal here is to map the relationship between AI and the SW, crystallize a cohesive research agenda, and raise several critical research issues. In particular, two early AI research efforts bear a striking resemblance to the SW initiative: Carl Hewitt's Open Information Systems Semantics[2] and Les Gasser's fundamental principles of distributed AI.[3]

What is the Semantic Web?

According to Tim Berners-Lee, "a semantic web is not artificial intelligence—the concept of machine-understandable documents does not imply some magical artificial intelligence which allows machines to comprehend human mumblings."[4] By juxtaposing "magical" and AI, Berners-Lee does the AI community an injustice, for the discipline of AI provides the most fundamental building blocks of the SW—and it is far from magical.

Fortunately, Berners-Lee provides a more constructive definition when he describes the SW as "an extension of the current one [Web], in which information is given well-defined meaning, better enabling computers and people to work in cooperation."[5] We can extend this definition to cover the general use of self-describing machine-readable knowledge that is accessible using standard Web programming constructs. Once knowledge sources have explicit and defined semantics, connecting them across the Web is what adds the semantic layer.

The SW is meant to enable an environment in which independent, Internet-connected information systems can exchange knowledge and action specifications, resulting in the execution of an activity acceptable to all systems involved. If designers built systems with this in mind originally, there might at least be some clear mapping between each system's data schemas. However, once we are dealing with heterogeneous information systems that have been developed independently,[6] little common ground exists on which to base such interactions. The SW's goal is to provide such enabling technologies. (See the "Semantic Web Data Representation" sidebar for more details.)

1094-7167/03/$17.00 © 2003 IEEE
Published by the IEEE Computer Society

History of Semantic Reasoning Systems

AI's history of research and development has created the foundations upon which the Semantic Web is being built. In the 1970s, knowledge representation and reasoning initiatives such as the Knowledge Representation Language, frame-based reasoning, and expert systems took center stage. The 1980s saw expert systems initiatives as the focus of Japan's Fifth Generation Project[1]; a drive to commercialize expert systems; and the development of new techniques such as the use of metaknowledge as a unifying element for integrating diverse knowledge. In the 1990s, intelligent-agent frameworks took the prior two decades' work into distributed multiagent environments. Semantic representation and reasoning, knowledge representation, frame-based representation, metalevel reasoning and abstraction, and multiagent systems have all been at the center of AI research at various points in time. Here are some of the most notable initiatives and related research, but this discussion is illustrative rather than exhaustive.

One of the earliest discussions of a machine-based system to create semantic agreement is attributable to Yehoshua Bar-Hillel in his quest for "fully automatic high quality translation." John Hutchins said of Bar-Hillel:

> With remarkable prescience, he realized that MT [machine translation] was an instance of "a well-known situation where accuracy may be traded for speed, and vice versa." For Bar-Hillel it was already "obvious" that "fully automatic MT, i.e. one without human intervention … [was] achievable only at the price of inaccuracy." The major obstacle to fully automatic translation was that there were no obvious methods "by which the machine would eliminate semantical ambiguities.[2]

The earliest attempt to represent knowledge using semantics can be traced back to Ross Quillian, who proposed the use of semantic networks for language understanding.[3,4] Groundbreaking work by Marvin Minsky in semantic information processing followed.[5] In 1975, Minsky suggested frame-based knowledge representation.[6] Current researchers are building methods for knowledge representation on the Semantic Web based on Minsky's foundations, one example of which is work by Ora Lassila and Deborah McGuinness.[7]

Research in distributed cooperating systems began with Alan Newell's work on the use of blackboard systems.[8] He envisioned an environment in which systems could deal with different representations at different levels of abstractions and communicate about them through a shared memory.

The development of KRL (Knowledge Representation Language) in 1977 was the first of what would become a stream of formal knowledge representation languages.[9] Later on, KIF (Knowledge Interchange Format)[10] and KQML (Knowledge Query and Manipulation Language)[11] would start to address interoperability issues for knowledge representation between heterogeneous knowledge sources.

Metalevel architectures and the use of metaknowledge to add a semantic layer began to generate a following in the 1980s. Michael Genesereth's work on metalevel architectures[12] and the collection by Pattie Maes and Danielle Nardi[13] presented many of the key foundations for semantic interoperability and cooperation. The metalevel approach has since been applied to areas ranging from blackboard system architectures[14] to knowledge acquisition and sharing.[15,16]

Combining the elements of knowledge representation, semantic agreement, and distributed reasoning became the mainstay of research into multiagent systems and distributed AI in the 1990s. Work on systems such as the Open Agent Architecture,[17,18] Flipside,[19] and InfoSleuth[20,21] continues.

References

1. E.A. Feigenbaum and P. McCorduck, *The Fifth Generation: Artificial Intelligence and Japan's Computer Challenge to the World*, Addison-Wesley, 1983.
2. W.J. Hutchins, *Milestones in Machine Translation*, http://ourworld.compuserve.com/homepages/WJHutchins/Miles-3.htm.
3. M.R. Quillian, "A Notation for Representing Conceptual Information: An Application to Semantics and Mechanical English Paraphrasing," Dept. of Energy tech. report SP-1395, System Development Corp., Santa Monica, Calif., Oct. 1963.
4. M.R. Quillian, "Semantic Memory," *Semantic Information Processing*, M. Minsky, ed., MIT Press, 1968, pp. 227–270.
5. M. Minsky, ed., *Semantic Information Processing*, MIT Press, 1968.
6. M. Minsky, "A Framework for Representing Knowledge," *The Psychology of Computer Vision*, P. Winston, ed., McGraw-Hill, 1975, pp. 211–277.
7. O. Lassila and D. McGuinness, "The Role of Frame-Based Representation on the Semantic Web," *Linköping Electronic Articles in Computer and Information Science*, vol. 6, no. 5, 2001, www.ep.liu.se/ea/cis/2001/005.
8. A. Newell, "Some Problems of Basic Organization," *Self-Organizing Systems*, M. Yovitz, G. Jacobi, and G. Goldstein, eds., Spartan, 1962, pp. 393–423.
9. D.G. Bobrow and T. Winograd, "An Overview of KRL, a Knowledge Representation Language," *Cognitive Science*, vol. 1, 1977, pp. 3–46.
10. R. Genesereth and R. Fikes, *Knowledge Interchange Format*, tech. report Logic-92-1, Computer Science Dept., Stanford Univ., 1992.
11. T. Finin and G. Wiederhold, *An Overview of KQML: A Knowledge Query and Manipulation Language*, tech. report, Computer Science Dept., Stanford Univ., 1991.
12. M.R. Genesereth, "An Overview of Metalevel Architecture," *Proc. Third Nat'l Conf. Artificial Intelligence* (AAAI 83), Morgan Kaufmann, 1983, pp. 119–123.
13. P. Maes and D. Nardi, *Meta-Level Architectures and Reflection*, Elsevier, 1988.
14. D.G. Schwartz and L. Sterling, "Using a Prolog Meta-Programming Approach for a Blackboard Application," *ACM Applied Computing Rev.*, vol. 1, no. 1, Jan. 1993, pp. 26–34.
15. T.J. Menzies, "Towards Situated Knowledge Acquisition," *Int'l J. Human-Computer Studies,* vol. 49, no. 6, 1998, pp. 867–893.
16. T. Menzies et al., "Issues with Meta-Knowledge," *Int'l J. Software Eng. and Knowledge Eng.*, vol. 10, no. 4, Aug. 2000, pp. 549–555.
17. P.R. Cohen et al., "An Open Agent Architecture," *Proc. AAAI Spring Symp.*, AAAI Press, 1994, pp. 1–8.
18. D. Martin, A. Cheyer, and D. Moran, "The Open Agent Architecture: A Framework for Building Distributed Software Systems," *Applied Artificial Intelligence*, vol. 13, nos. 1/2, 1999, pp. 91–128.
19. D.G. Schwartz, *Cooperating Heterogeneous Systems*, Kluwer Academic Publishers, 1995.
20. D. Woelk, M. Huhns, and C. Tomlinson, "InfoSleuth Agents: The Next Generation of Active Objects," *Object Magazine*, July/Aug. 1995, pp. 33–40.
21. R. Bayardo et al., "Semantic Integration of Information in Open and Dynamic Environments," *Proc. ACM SIGMOD Int'l Conf. Management of Data*, ACM Press, 1997, pp. 195–206.

Semantic Web Data Representation

We can approach data representation in a Semantic Web environment as a series of layers, each of which adds a degree of semantic depth to a given data model. Figure A shows the four basic layers: XML, RDF (Resource Description Framework), Ontology, and Logic.

XML Schema

The XML Schema standard forms a broad base on which developers can build interoperable XML applications. It enables the cross-organizational sharing and verification of documents. The schema specification consists of two parts: a language to describe the high-level structure of an XML document, and a list of allowable data types that can be used in those documents.

Using XML Schema lets an XML parser identify not only the syntax of XML data but also the correctness and completeness of that data using its structure. This enables the easy identification of missing data, data-formatting errors, and out-of-range values. Some major ontology initiatives have adopted an XML-based semantic markup language for expressing standard semantics on the Web.[1] Both DAML (DARPA Agent Markup Language)[2] and OIL (Ontology Inference Layer)[3] use forms of XML as the basis for their ontological representations. The following example shows a simple XML representation of two common terms: bird and cat.

```
<?xml version="1.0">
<birds>
    Animals that fly:
    <bird>
        <name>Tweety</name>
        <color>Yellow</color>
        <address>Granny's Cage</address>
        <eats>birdseed</eats>
        <enemy>puddy tat</enemy>
    </bird>
</birds>
<cats>
    Animals that eat birds:
    <synonym>puddy tat</synonym>
    <cat>
        <name>Sylvester</name>
        <color>gray</color>
        <address>Granny's House</address>
        <eats>birds</eats>
    </cat>
</cats>
```

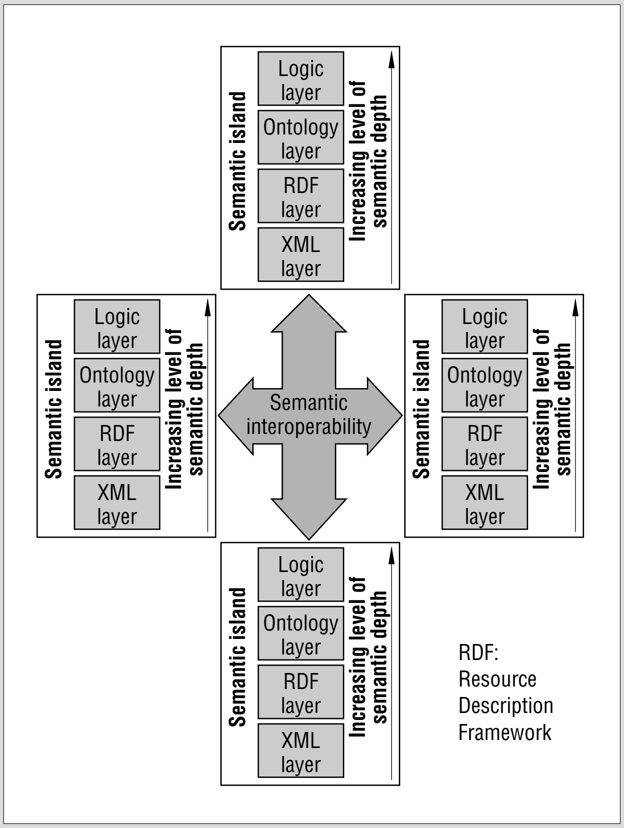

Figure A. Seeking interoperability across multiple layers of increasing semantic depth.

XML Schema stops short, however, at allowing for true semantic representation. The Semantic Web initiative advocates the addition of RDF on top of XML to begin dealing with that task.

Resource Description Framework

RDF is essentially a rudimentary, Web-enabled language of Subject, Verb, Object triples that we can use to represent relationships between data. In that sense, it provides Web-based applications with simple semantic-net-like capabilities much like the venerable (Tweety) isa (bird) and (bird) has (wings) representations. The key dif-

Open Information Systems Semantics

Hewitt proposed OISS as a response to the increasing complexity engendered by the use of large-scale open systems:

> OIS Semantics provides an account of the meaning of Representational Activity (i.e. conveying information about System Commitments that is accomplished using digital communications). Representational Activities entail changes in Systems Commitments, and that change is the Meaning of the activities. This is an open-world characterization of meaning—as opposed to previous close[d]-world attempts based on possible worlds in which the meaning of a set of sentences is defined to be the set of all possible worlds that satisfy the meaning conditions. In logical semantics, representation is the mapping between a sentence proposition and specified meaning conditions. Meaning is built on and grows out of representation. Two participants agree about the meaning of a sentence when they agree about the meaning conditions.[2]

An *open system* is any system that is subject to unanticipated outcomes in its operation and can receive new information from outside itself at any time. Most classic information systems do not fall into this category. Although an information system might receive input from the outside at any time, these inputs are generally anticipated and match predefined syntactic structures and schemas. Distributed AI enabled, for the first time, nondeterministic communication between two or more indeference here is that the elements of RDF statements can be uniform resource identifiers (URIs), so that their location can be referenced anywhere on the Web.

More in-depth discussion of the roles of XML and RDF in the SW is available elsewhere.[4,5] Table A shows a brief example of three fictitious RDF triples that theoretically could be used in reasoning about Tweety.

Table A. RDF descriptions of Tweety.

Subject	Verb	Object
www.tweetybird.xrg	created_by	www.warnerbros.xom/legal
www.tweetybird.xrg/tweety	isa	www.myOntology.xrg/bird
www.myOntology.xrg/bird	has_part	www.ScienceOntology.xrg/life/wings

Ontology

The use of ontology is fundamental to a viable SW.[6] Without it, at best we end up with semantic islands—not much use for interoperability and communication.

Currently, we can use the XSLT (Extensible Stylesheet Language Transformations) standard to map multiple XML data representations, but it is a difficult, manual task. An alternative is to use a central ontology as a touchstone from which and to which multiple XML representations are mapped. Mapping between XML representations is a core SW function, one made possible through the use of ontology technologies.

Did you ever notice what happens if you try to run the plural form of "ontology" through your spell checker? While the term "ontologies" sounds right (and looks right), there is, in fact, no such animal. For by its very definition, "ontology" connotes a unified view—a single universally agreed-upon system of meanings. If you have more than one ontology, you have no ontology. So, to get RDF "islands" to communicate, there must be a common ontology that connects the terms in one to the terms in the other. Of course, achieving a single common ontology is not realistic, so the actual goal must be effective mechanisms to integrate and reconcile overlapping and inconsistent ontology islands—hence the centrality of the semantic interoperability function as Figure A shows.

Mapping XML through ontology is therefore just the first step, as there is no way to guarantee a given ontology's completeness. Therefore, we must address the mapping of multiple "ontologies" in much the same way as we address multiple XML and RDF representations. An ontological mapping is a means of achieving interoperability between ontologies. This, of course, ignores our earlier proclamations of a single ontology. However, in practice multiple partial ontologies will exist, and reconciling them is essential to a viable SW. For more on this, see the "Additional Research Directions" sidebar.

Logic

The use of first-order or higher-order logic is also fundamental to a viable SW. Without it, at best we end up with a network of concepts with an agreed-upon grounding but no inferential abilities. First-order logic allows us to reason across the RDF elements defined in a SW environment. Stated in terms of the earlier XML/RDF example, if we use RDF to express www.tweetybird.xrg/tweety isa www.myOntology.xrg/bird and www.dangerouscats.xrg/sylvester isa www.myOntology.xrg/cat, we then need a logic that lets us represent simple rules such as

A enemy B <— A isa www.myOntology.xrg/bird,
B isa www.myOntology.xrg/cat.

Further reading on SW layers and formalization is available:

- Semantic Web initiative: www.w3.org/2001/sw
- XML specification: www.w3.org/XML
- RDF specification: www.w3.org/RDF
- Requirements for a Web Ontology Language: www.w3.org/TR/2002/WD-webont-req-20020708
- Use of logic for the SW: www.w3.org/DesignIssues/Logic.html

References

1. B.R. Gaines and M.L.G. Shaw, "Knowledge Acquisition, Modelling and Inference through the World Wide Web," *Int'l J. Human-Computer Studies*, vol. 46, no. 6, June 1997, pp. 729–759.
2. J. Hendler and D.L. McGuinness, "The DARPA Agent Markup Language," *IEEE Intelligent Systems*, vol. 15, no. 6, Nov./Dec. 2000, pp. 72–73.
3. D. Fensel et al., "OIL in a Nutshell: Knowledge Acquisition, Modeling, and Management," *Proc. European Knowledge Acquisition Conf.* (EKAW 2000), LNAI 1937, R. Dieng et al., eds., Springer-Verlag, 2000, pp. 4.1–4.12, www.cs.vu.nl/~ontoknow/oil/downl/oilnutshell.pdf.
4. S. Decker et al., "The Semantic Web—The Roles of XML and RDF," *IEEE Internet Computing*, vol. 4, no. 4, Sept./Oct. 2000, pp. 3–13.
5. D. Fensel, "The Semantic Web and Its Languages," *IEEE Intelligent Systems*, vol. 15, no. 6, Nov./Dec. 2000, pp. 67–73.
6. T. Gruber, "A Translation Approach to Portable Ontology Specifications," *Knowledge Acquisition*, vol. 5, no. 2, 1993, pp. 199–220.

pendent systems, resulting in the exchange of inputs that had not been predefined—often referred to as the *open system* or *open world* assumption. This realization led to a series of nontrivial questions regarding the nature of interaction between systems.

Has anything really changed?

When Hewitt presented his far-reaching and insightful analysis of OISS,[2] it appeared alongside Gasser's equally insightful view of what does and does not change in AI when you embrace the open system assumption.[3] When these works appeared, they were not "initial thoughts" but rather the analytical culmination of research in distributed AI that had begun two decades earlier.

Together, the open system assumption and the principles of DAI comprise the single most important leap in AI's development from a niche research area to the science that is now forming the SW core. Much before the SW concept became popular, the opposing and slightly controversial approaches had already laid the groundwork for the main SW issues being dealt with today. Consider the following quotations:

> This is an open-world characterization of meaning—as opposed to previous closed-world attempts.[2] (Hewitt, 1991)
>
> One consistent difference between the Semantic Web and many data models of programming languages is the "closed world assumption."[4] (Berners-Lee, 1998)

Identifying the open-world assumption was a crucial turning point in the development of both OISS and the SW. This commonality between the two does not end here; rather, it is just the point of embarkation on a common path.

Hewitt makes four fundamental assertions that are important in understanding OISS:

- *Negotiations* play a fundamental role.
- *Robustness*, in terms of the ability to keep commitments in the face of conflict, must be supported.
- *Scalability*, in terms of the ability to increase the scale of commitments, is crucial.
- *Meaning* is derived from a representational activity that is the resulting change in that activity. Without communication, there is no representational activity, where representational activity involves communicating information regarding system commitments. Or, in Hewitt's words (p. 94), "No Representation without communication!"[2]

Seven core concepts

Hewitt further identifies seven concepts that are core to the success of OISS. They are equally essential for the SW to succeed as an enabling infrastructure. By considering each one from the OISS perspective, we can begin to understand how we must adopt and adapt these same elements for the SW.

Conflicts. The ability to detect and resolve conflicts is fundamental in any open-systems interaction and in a SW environment. OISS conflicts result from the asynchrony of multiple participants operating independently. SW conflicts are possible at multiple levels, from conflicting XML and RDF (Resource Description Framework) representations to incompatible ontology terms. They lead to incomplete or contradictory first-order-logic representations deployed in support activities. A SW implementation that does not include a mechanism to resolve conflicts will be ineffective once communication moves beyond a controlled environment.

Trials of strength. These occur whenever a conflict must be resolved; they are common especially when ontologies conflict (see the "Additional Research Directions" sidebar). Hewitt proposes mechanisms for processing OIS trials of strength,[2] many of which apply to the SW.

System commitments. These are the basis for joint activities: they let multiple participants agree to perform a joint processing activity. System commitments involve allocation of resources, agreement on processing logic, and agreement on data semantics. The importance of joint processing in the SW is apparent even if we limit ourselves to data interoperability. For example, systems participating in a consortium might commit to using the same base ontology for a subset of their representational requirements. They can further agree that if that ontology becomes unavailable within a committed time constraint, they will use an agreed-upon alternative.

Deduction. Deduction is based on a given system's internal logical structure. The SW initiative views this as a role to be filled by a Logical layer above RDF representations. It does not present any particularly difficult research problem, but is another example of the common ground being shared. Considering that having multiple mechanisms for deduction can lead to conflict in both OISS and the SW, a cooperation mechanism is necessary.

Cooperation. When the participating systems have mutually dependent roles in system commitments, they need to cooperate to establish those roles. For example, they might have established the commitment to an alternate ontology described earlier through such a mechanism. Likewise, when participants have multiple deductive mechanisms available, they must cooperate to navigate the interaction between them.

Representational activity. This is any OISS activity that results in a change in system commitments as defined earlier. The meaning of a representational activity in OISS can be localized; it can also propagate out to other participants mitigated by trials of strength and cooperation. Similar behavior can be expected in a SW environment, as meaning will need to be propagated from one uniform resource identifier (URI) to another.

Negotiation. This is a modified trial of strength in which the participants communicate various representational activities that change system commitments. Negotiating mechanisms can be complex, and they form a rather active AI subdiscipline. It will be interesting to see how much of AI's contribution to automated negotiation will be applied to the SW environment, as meaning will need to be mapped from one URI to another.

Adding social conceptions of knowledge and action

Gasser's learned analysis of OISS presents a further refinement of the key issues related to understanding and problem solving in open systems.[3] His six principles form the basis of distributed AI:

- AI research must set its foundations in ways that treat the existence and interaction of multiple actors as a fundamental category.
- Tension will exist between situated pragmatic knowledge and action-at-a-distance in alternative contexts.
- DAI representation and reasoning approaches must assume that multiple representations are recursively possible at any level of analysis or action and that actors will employ multiple representations individually and collectively. DAI approaches must also provide mechanisms for reasoning among multiple representations.
- DAI theory and practice must account for resource-limited activity.
- DAI theory and practice must provide accounts of and mechanisms for handling three key problems: joint qualification, representation incommensurability, and failure indeterminacy.
- DAI theory and practice must account for how groups of agents can achieve joint, robust, and ongoing courses of action despite indeterminate foul-ups and inconsistencies, which can occur recursively at any level of the system.

Aside from the first principle, which is more philosophical in nature, the distance between Gasser's guiding principles and the demands of the current SW initiative is short indeed. Designers of SW technologies must expect tension between situated pragmatic knowledge at a given URI and alternative representations in other locations—that is the SW's essence. A robust SW implementation cannot avoid handling recursive representation and reasoning, even if current RDF research does not yet deal with this. Likewise, can we have a SW that does not consider bounded resources and multiple conflicting representations? No, these are fundamental. Also, some SW interactions will necessitate joint courses of action determined by the participants. Anything less misses the vision. In fact, everything in Hewitt's OISS and Gasser's DAI principles must at least be considered if not fully addressed to fulfill the SW vision.

Additional Research Directions

In the Deductive Indecision Problem, two serial operations—for example, two withdrawals from the same account by different participants—occur such that it is impossible to physically determine which operation will be executed given the account's balance limits. If an account with a balance of 100 exists on System A, and two valid messages ("Withdraw 70" and "Withdraw 80") are concurrently sent to System A, Carl Hewitt says, "No amount of knowledge of the physical circumstances in which the withdrawal requests are made determines the outcome. Therefore, the outcome cannot be *deductively decided* even from complete knowledge of all circumstances" at the time both requests were made.[1]

This problem is well-known in the database management literature yet must be revisited in terms of an open Semantic Web environment. Semantic interoperability at the data level, as Hewitt describes, is not enough. Inference related to Web participants' actions must be able to deal with the level of indecision that the Deductive Indecision Problem addresses and still continue functioning. In an environment in which long duration and indecisive transactions are the norm rather than the exception, new mechanisms must be put into place. One attempt to address part of this problem is Alfred Loo and Y.K. Choi's distributed selection algorithm for the Internet.[2]

One of the motivations behind using ontologies is the desire to set a semantic standard that many unrelated participants can share.[3] This goal is sabotaged as soon as multiple ontologies appear in the same conceptual space—an inevitable phenomenon. So, establishing some notion of semantic equivalence—when Concept A in Ontology X has the same meaning as Concept B in Ontology Y—is important.[4,5]

The classic example is that of a ZIP code.[6] Zip in Ontology A refers to a string of digits that we can use to establish an entity's physical location in the world. We might refer to the identical concept in Ontology B as Postal Code. How do we establish an equivalence relationship between the two ontologies? Berners-Lee and his colleagues[6] suggest that the ontologies themselves provide an equivalence relationship stating that one's postal code is equivalent to the other's ZIP code. But taking that approach begs the question by expecting a situation in which such equivalence knowledge is available to the creators of one or both ontologies.

A more practical albeit significantly more difficult approach would be to establish some mechanism to compare ontologies and infer the equivalence of concepts based on similarity measures. Alexander Maedche and Steffen Staab propose some useful similarity measures.[7] Peter Weinstein and William Birmingham propose a methodology for determining the compatibility of multiple ontologies.[8]

Other areas of research that directly affect the SW's future include cooperative multiagent systems[9] and Semantic Web languages.[10]

References

1. C. Hewitt, "Open Information Systems Semantics for Distributed Artificial Intelligence," *Artificial Intelligence*, vol. 47, no. 1–3, 1991, p. 85.
2. A. Loo and Y.K. Choi, "A Peer-to-Peer Distributed Selection Algorithm for the Internet," *Internet Research: Electronic Networking Applications and Policy*, vol. 12, no. 1, 2002, pp. 16–30.
3. D. O'Leary, "Using AI in Knowledge Management: Knowledge Bases and Ontologies," *IEEE Intelligent Systems*, vol. 13, no. 3, May/June 1998, pp. 34–39.
4. D.G. Schwartz, "When Email Meets Organizational Memories," *Int'l J. Human-Computer Studies*, vol. 51, no. 3, 1999, pp. 599–614.
5. D.G. Schwartz, "Shared Semantics and the Use of Organizational Memories for Email Communications," *J. Internet Research*, vol. 8, no. 5, 1998, pp. 434–441.
6. T. Berners-Lee, J. Hendler, and O. Lassila, "The Semantic Web," *Scientific Am.*, May 2001.
7. A. Maedche and S. Staab, "Measuring Similarity between Ontologies," *Proc. 13th European Conf. Knowledge Eng. and Knowledge Management* (EKAW), LNAI 2473, Springer-Verlag, 2002.
8. P. Weinstein and W. Birmingham, "Comparing Concepts of Differentiated Ontologies," *Proc. 12th Workshop Knowledge Acquisition, Modeling, and Management* (KAW 99), 1999, http://sern.ucalgary.ca/KSI/KAW/KAW99/papers/Bermingham1/kaw99_final.pdf.
9. V.R. Lesser, "Cooperative Multiagent Systems: A Personal View of the State of the Art," *IEEE Trans. Knowledge and Data Eng.*, vol. 11, no.1, Jan./Feb. 1999, pp. 133–142.
10. I. Cruz et al., *The Emerging Semantic Web*, IOS Press, 2002.

The Semantic Web research agenda

The principles of DAI and OISS represent rich streams of research that embody AI's contribution to the next generation of the Web.

Clearly, the seven points of OISS and six principles of DAI provide a fairly robust research agenda for the SW. In addition, several research issues have direct implications for the SW initiative but are not addressed at the core of OISS. These include questions dealing with the Deductive Indecision Problem (which Hewitt also discusses[2]) and the comparison of diverse ontologies. There are other such starting points for effective SW research that provide clear directions as to the real issues that must be addressed. One additional example is Victor Lesser's work,[7] which, like that of Hewitt and Gasser, identifies key principles to be used in building multiagent systems. At the end of the day, that's what the Semantic Web is poised to become: an immense multiagent system. (For more details, see the "Additional Research Directions" sidebar.)

What's been accomplished so far?

Many SW research groups have emerged over the past four years. The most well-known are perhaps DARPA's Agent Markup Language (DAML) project, The OntoAgents and Scalable Knowledge Composition groups at Stanford University, and the Ontobroker group at the University of Karlsruhe—all of which have a heavy ontology and knowledge representation focus.

There is no lack of "Semantic Web" applications that predate the SW. Consider, for example, the Unified Medical Language System,[8] which uses a semantic network to provide a consistent categorization of all concepts found in a medical metathesaurus and provides links between those concepts at the semantic-type level. The three main UMLS knowledge sources provide powerful tools for enhancing medical-information retrieval.

Another pre-SW initiative, the venerable and ambitious Cyc project,[9] has recently reclothed in SW garb. It is perhaps indicative that Cyc, which started as a research project in 1984 and has been applied to numerous application domains along the way including the integration of heterogeneous databases, is now being reincarnated as a tool for—or rather a potential piece of—the SW.

There are already initial successes in applying SW technologies to areas such as search and data interoperability. Search can easily benefit from the addition of a semantic layer, so it has become the early focus of much SW research. In today's environment, *search* is popularly found in juxtaposition to *Web*, since the Web presents a scale and complexity of search problem heretofore unseen—even in the heuristic-search genre that formed the core of early AI research. Given a search term, semantic search uses a combination of knowledge base references and associations to augment search results based on semantic relationships. The Ontology Web Language is one such effort that is seeing early success.

Data interoperability, while not limited to Web environments, is a point of great pain in many commercial computing environments. So, this area too has begun to develop a significant level of research aimed at adding a semantic layer to create interoperable data architectures. The goal here is to help reconcile different representations in large multidatabase computing environments. Adding a semantic layer results in lower maintenance costs and fewer development errors, as developers can more easily identify both conflicting and complementary data items once they establish their semantic equivalence. Much of this work is being done with the SW as its guide—in terms of both the underlying XML/RDF technology being used and future architectural goals leading to interorganizational data interoperability.

Some individual pieces of the SW puzzle are already being adopted. For example, the US electric power industry is using RDF to exchange power system models between system operators. The Electric Power Research Institute Common Information Model is meant to serve the power industry in much the same way UMLS serves the medical profession. In fact, the North American Electric Reliability Council has mandated that the XML-based framework be the standard for exchanging models between transmission system operators.

There are, at last count, over 21 European Information Society Technologies projects on semantics in distributed systems. They range from the Corporate Ontology Grid (www.cogproject.org), which uses ontological modeling to integrate corporate information into a commercial interorganizational grid, to the WonderWeb Ontology Infrastructure for the Semantic Web project (http://wonderweb.semanticweb.org), which is working to establish a Web standard ontology language and ontological engineering tools.

The success of these applications tells us that there is immense potential in the SW initiative. This potential derives its strength from the Internet's inherent distribution and reach, enabling a SW that is in fact a worthy evolutionary step down the road of AI. These early adoptions and new projects have one crucial thing in common. Successful though they might be, they focus almost exclusively on knowledge representation, annotation, tagging, retrieval, and interpretation. They've not yet taken the leap to a real heterogeneous distributed-agent environment. When they do, they will quickly discover that the issues identified by OISS and DAI are alive and well and lying in wait.

The Semantic Web initiative should not, and will not, develop as an attempt to displace the wealth of existing AI tools and techniques. On the contrary, the universal connectivity, addressing, and transport mechanisms offered by combining TCP/IP, URIs, and RDF/XML wrappers might just provide the right glue to combine with Open Knowledge Base Connectivity, the Knowledge Query and Manipulation Language, or other such attempts to combine reasoning with ontology-based knowledge representation.

As Robert Frost so eloquently wrote,

> I shall be telling this with a sigh
> Somewhere ages and ages hence:
> Two roads diverged in a wood, and I—
> I took the one less traveled by,
> And that has made all the difference.

The road we take does make all the difference but, contrary to Frost's experience, we would do well to consider the road already traveled. If we are to create the next generation of SW infrastructure, we must pick up where Hewitt, Gasser, and other AI pioneers left off and address the issues of conflicts, trials of strength, system commitments, deduction, cooperation, representational activity, and negotiations. It makes just as much sense when it comes to the SW as it did with respect to OISS. Anything less, and the vision of the SW is doomed to failure—even if we don't call it AI.

References

1. T. Berners-Lee et al., "World-Wide Web: The Information Universe," *Internet Research: Electronic Networking, Applications and Policy*, vol. 2, no. 1, 1992, pp. 52–58.
2. C. Hewitt, "Open Information Systems Semantics for Distributed Artificial Intelligence," *Artificial Intelligence*, vol. 47, nos. 1–3, 1991, pp. 79–106.
3. L. Gasser, "Social Conceptions of Knowledge and Action: DAI Foundations and Open System Semantics," *Artificial Intelligence*, vol. 47, nos. 1–3, 1991, pp. 107–138.
4. T. Berners-Lee, *Web Design Issues: What the Semantic Web Can Represent*, 1998, www.w3.org/DesignIssues/RDFnot.html.
5. T. Berners-Lee, J. Hendler, and O. Lassila, "The Semantic Web," *Scientific Am.*, May 2001, www.scientificamerican.com/article.cfm?articleID=00048144-10D2-1C70-84A9809EC588EF21&catID=2.
6. D.G. Schwartz, *Cooperating Heterogeneous Systems*, Kluwer Academic Publishers, 1995.
7. V.R. Lesser, "Cooperative Multiagent Systems: A Personal View of the State of the Art," *IEEE Trans. Knowledge and Data Eng.*, vol. 11, no. 1, Jan./Feb. 1999, pp. 133–142.
8. C.R. Selden and B.L. Humphreys, *Unified Medical Language System (UMLS): January 1986 through December 1996: 280 Citations*, Nat'l Library of Medicine, 1997; www.nlm.nih.gov/pubs/cbm/umlscbm.html.
9. R.V. Guha and D.B. Lenat, "Cyc: A Midterm Report," *AI Magazine,* vol. 11, no. 3, 1990, pp. 32–59.

The Author

David G. Schwartz is a senior lecturer and former head of the information technology division at the Graduate School of Business Administration of Bar-Ilan University. He is also a founding partner of Apropos IT Ventures, a venture capital fund focusing on early-stage IT companies in Israel. He is editor of the *Journal of Internet Research.* He received his PhD in computer science from Case Western Reserve University. Contact him at the Graduate School of Business Administration, Bar-Ilan Univ., Ramat-Gan, Israel 93220; dschwar@mail.biu.ac.il.