

Preparing for the Semantic Web

Antonio M. Lopez, Jr.
Computer Sciences and Computer Engineering Department
Xavier University of Louisiana
New Orleans, LA 70125, USA

Abstract

As today's undergraduates in computer information systems work towards learning those skills needed in the world of e-Commerce (e.g., HTML, Java, XML, etc.), an enhancement to the World Wide Web is being "spun." The Semantic Web is the envisioned end-state for the movement of the World Wide Web from words, images, and audio understood only by humans to those same things "wrapped in" organizing concepts and relationships understood by both humans and software agents. Since the Semantic Web is a research effort under the auspices of the World Wide Web Consortium, how are faculty members to expose students to the emerging technology that will impact how e-Commerce will be supported in the future?

This paper highlights some of the World Wide Web Consortium work done thus far in moving the current Web toward a Semantic Web. It next presents an overview of ontology development, the key enabling technology for the Semantic Web. Finally, it shows how undergraduate researchers, mentored by their faculty members, can develop ontologies that lead naturally into preparing for the Semantic Web.

Keywords: World Wide Web, e-Commerce, semantic web, ontology, undergraduate research

1. INTRODUCTION

The World Wide Web (WWW) is a popular communications medium for the exchange of information among people – i.e., individuals or people who are members of or represent organized groups (non-profit) or businesses (profit). The computers that people use, whether sitting on their desks at work, in their homes, accompanying them on business trips, in cyber-cafes where they socialize, or in public libraries where they study, are all potential points of entry into the WWW network of information exchange and business transactions. Unfortunately, the success of the communications medium and the large volume of information that it provides have actually hampered e-Commerce growth (Davies et al. 2003), both e-Business (B2B) and Web-Commerce (B2C). In B2B for example, today's businesses are coming to the realization that their knowledge is a

valuable corporate asset, and this knowledge needs to be shared with global partners so that each can improve their competitive edge and increase their market shares. Competitive businesses must find effective ways of working together to achieve their goals. RosettaNet (www.rosettanet.org) is an example of a non-profit consortium of major information technology, electronic components, and semiconductor manufacturing businesses working to create and implement industry-wide, open business process standards. In B2C, today's customers usually spend a significant amount of time browsing for businesses that sell a desired product and, typically, only a few such businesses are browsed before customer exhaustion occurs and the customer settles for what has been found. Competitive businesses must find effective ways of reaching customers. Makebuzz (www.makebuzz.com) is an example of a for

profit company that combines a massive bank of online marketing knowledge and artificial intelligence (AI) strategies to increase the chances that their client's products will be visible enough on a popular search engine listing so that a potential customer might select them for further product browsing.

While RosettaNet and Makebuzz are respectively today's solution to the corporation's *sharing of business knowledge problem* and the customer's *information overload problem*, tomorrow's solution can be found, in part, in the use of software agents (also called intelligent agents) that "understand" Web pages. Software agents are computer programs that work without direct human control or constant supervision to achieve goals set forth by humans (Stojkovic and Lupton 2000). If the WWW was understandable to software agents, much more so than today's keyword match or heuristic search of hyperlinks, then software agents could collect, filter, and process desired information found on the WWW for customers or businesses. At present, both elements of e-Commerce (B2B and B2C) are being held back by the lack of standards for (1) representation of the information contained on a Web page, (2) translation of that information into other more useful forms, and (3) description of content that is software agent understandable (Fensel 2001). If standards in these three areas are defined, then commercial software can be developed to automatically encode the semantics of a Web page as the Web page is being created. Furthermore, these standards will facilitate the creation of software agents that collect Web content from different sources, refine the content, and exchange it with other software agents.

2. THE SEMANTIC WEB

The Semantic Web was initially described as "the conceptual structuring of the Web in an explicit machine-readable way" (Berners-Lee and Fischetti 1999). Later, Berners-Lee et al. (2001) wrote, "The Semantic Web will 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." In an issue of *Business Week* (Port 2002), Richard Hayes-Roth, the chief tech-

nology officer for software at Hewlett-Packard, is quoted as saying, "We expect the Semantic Web to be as big a revolution as the original Web itself."

Starting with the current Web, a community of users must decide how information on a Web page can be given well-defined semantics, thus making this new Semantic Web page understandable to a software agent. The community of users that has accepted this challenge is the World Wide Web Consortium (W3C). W3C (www.w3c.org) is a very active organization having offices in Valbonne, France; Cambridge, MA, USA; and Tokyo, Japan with a full-time staff of more than 60 and more than 500 member organizations (Cherry 2002). Slowly but surely, W3C has been moving the Web toward the Semantic Web end-state. Two important steps were the extensible markup language (XML) and the resource description framework (RDF).

Presently, XML allows creators of Web pages to produce and use their own markup tags (Boggs 2002). However, if other users of the Web know the meaning of the XML tags, then they too can write scripts that make use of those tags. Thus, a small step has been taken to represent the information contained on a Web page and make it understandable to a software agent. Unfortunately, XML does not provide standard data structures and terminologies to describe business processes and the exchange of product information (Fensel 2001). So as far as a software agent is concerned, the translation issue (mentioned in the previous section as needing standardization) has not been sufficiently advanced with XML; thus the need for RDF. The RDF data model, which is equivalent to the semantic network formalism, consists of resources, properties, and statements written using XML tags (Gomez-Perez 2002). Universal Resource Identifiers (URIs) can be used in identifying resources, properties, and statements. The RDF Schema (RDFS) provides a means of defining relationships between resources and properties. Hence, RDFS provides the basics for defining knowledge models that are similar to frame-based systems.

The Defense Advanced Research Projects Agency (DARPA) has also supported the development of the Semantic Web. In an ef-

fort to make Web content more accessible and understandable to agents (human and software), DARPA has funded research in languages, tools, infrastructure, and applications. Based upon XML and RDF, the DARPA Markup Language (DAML) was developed and coupled with the Ontology Inference Layer (OIL) to produce DAML + OIL, a proposed starting point for the W3C's Semantic Web activity's known as the Ontology Web Language (OWL) (McGuinness et al 2002). Web integration, frame-based systems, and description logics inspired the specification of OWL. W3C's most recent release of DAML + OIL was March 2001, and the deadline for the call for comment on the "last working-draft" of OWL was March 2003.

structuring the Web into the Semantic Web. Figure 1 is a semantic network with some basic knowledge that someone may want to impart on a Web page (e.g., This might be the beginnings of a logical theory: Men and women are humans.).




Figure 1.

We conclude this section with a small example that demonstrates research work in

```

<?xml version="1.0" encoding="iso-8859-1" ?>
<rdf:RDF
  xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
  xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
  xmlns:daml="http://www.daml.org/2001/03/daml+oil#" >

  <rdf:Description about="http://www.w3.org">
    <Publisher> World Wide Web Consortium </Publisher>
  </rdf:Description>

  <daml:Class rdf:ID="Human">
    <rdfs:label>Human</rdfs:label>
  </daml:Class>

  <daml:Class rdf:ID="Man">
    <rdfs:subClassOf rdf:resource="#Human"/>
  </daml:Class>

  <daml:Class rdf:ID="Woman">
    <rdfs:subClassOf rdf:resource="#Human"/>
  </daml:Class>

  <daml:Class rdf:about="#Human">
    <daml:disjointUnionOf rdf:parseType="daml:collection">
      <daml:Class rdf:about="#Man"/>
      <daml:Class rdf:about="#Woman"/>
    </daml:disjointUnionOf>
  </daml:Class>

  <daml:Class rdf:ID="HumanBeing">
    <daml:sameClassAs rdf:resource="#Human"/>
  </daml:Class>

</rdf>

```

Table 1

Table 1 shows the markup language implementation of the semantic network adding information about who might have published the Web page and two new ideas. The first idea is that every human is either a man or a woman (disjointness). The second idea is that the concept of a human is the same as the concept of a human being (sameness). The encoding using XML, RDF, RDFS, and DAML is straightforward but still subject to change as W3C continues its research work (i.e., OWL). However, the underpinning is without doubt the ontological structures represented, in part, by semantic networks.

3. ONTOLOGY DEVELOPMENT

Ontologies are a key enabling technology for the Semantic Web because of what ontologies promise – a shared and common understanding of a domain that can be communicated between people and software systems (Davies et al. 2003). An ontology is a logical theory which gives an explicit, partial account of a conceptualization; it is an intentional semantic structure that encodes the implicit rules constraining the structure of a piece of reality (Guarino and Giaretta 1995). In the early 1990s AI researchers struggled with concepts such as metadata (Lopez and Saacks 1992), meta-knowledge (Lopez 1993), and their combined use in information management systems (Saacks-Giguette and Lopez 1993) to provide a vocabulary of terms and relations that would facilitate the development of some domain model (The previous three references are rooted in a NASA domain called the Earth Observing System). Gruber (1993) was the first to use the term “ontology,” which he wrote was “a specification of a conceptualization.” Since then, the AI community has adopted ontology development as a pre-requisite to building knowledge-based systems because every knowledge model has an ontological commitment (Noy and Hafner 1997); that is, the ontology captures the set of concepts, terms, and relationships used to describe the knowledge of the software system.

Maedche (2002) gives a mathematically rigorous definition of an ontology structure, one that developers of the Semantic Web will rely upon to organize the underlying metadata and meta-knowledge of a domain for the purpose of comprehensive and transportable software agent understanding. In

sum, an ontology structure consists of two disjoint sets, a set of concepts and a set of relations. An ontology structure has a concept hierarchy or taxonomy expressing relations between concepts, for example “Woman *subClassOf* Human”, which is found in Figure 1 and encoded in Table 1. It also has function relations that relate concepts non-taxonomically (others would call these attributes of concepts with values to be established with specific instances of that concept). Finally, the ontology has axioms expressed in an appropriate logical language (e.g., first order logic). In Table 1, we embedded the ontology axiom “Human *disjointUnionOf* Man and Woman.” As will be seen below, for our purpose of preparing students for the Semantic Web, we take a more informal, graphical approach that can be encoded in a first order predicate calculus programming language. Our experience is that undergraduates easily understand this approach to ontology development.

Ontology development can be accomplished top-down (Lenat 1995) or bottom-up (van der Vet and Mars 1998). A top-down development of an ontology structure for a large knowledge-based system can take years to construct, so developers do well to import existing top-level ontologies and combine them with knowledge from several independently developed bottom-up ontologies. Ontologies developed bottom-up are called domain ontologies, and they contain the terms that are useful in a wide range of different applications within a specific domain. In a mature domain such as high school algebra word problems there is widespread agreement on the basic terms and relationships, for example objects (people, car, etc.), actions (walk, drive, etc.), algebraic relationships, and types of problems (distance-rate-time, mixture, work, etc.). However in a domain such as religion, there is not as much widespread agreement. Domain ontologies have been successfully developed for military applications (Valenti et al. 1999; Bowman et al. 2001) as well as library applications (Weinstein and Alloway 1997; Welty and Jenkins 1999).

Although differences can exist conceptually regarding what should or should not be in an ontological structure, there is some general agreement on a number of basic issues. As Chandrasekaran et al. (1999) explained: (1)

A domain ontology contains the names of objects that are found in the domain. (2) Relationships exist between the objects. (3) Objects can have parts. (4) Objects have attributes that can take on values. (5) Attributes and relationships can change over time. (6) There are processes that occur over time in which objects participate. To illustrate these points of agreement consider the Shoe domain ontology (Figure 2). The objects in the Shoe ontology are the nodes of the diagram (e.g., Moccasins, Boots, etc.). Labeled lines connecting nodes express relationships between objects (e.g., subClassOf, aPartOf). Attributes are rectangular tags attached to the node (e.g., size, color, brand). Attributes of a node are inherited by the subclasses of a node, and eventually attributes take on values for very specific instances of a concept (e.g., John's Court_shoes are size: 11W, color: white, and brand: Reebok). The temporal nature of the ontological structure cannot be shown in a diagram such as Figure 2; however, it is understood and must be coded (e.g., date-time stamp) when the ontological structure is implemented so that a software agent can

use it. The software agent must understand that John did not always have size 11W court shoes; he grew into them over time. Perhaps, in that time period, there is a record that John regularly purchased court shoes.

To be useful, a domain ontology must be a highly internally coherent logical theory having specific "gates" through which interaction with other domain ontologies can occur. The question is: How does a developer test the coherence of the logical theory? Fortunately, the DAML + OIL specification includes both a first-order logic semantics and a model-theoretic semantics. These semantics enable the use of a first order predicate calculus programming language such as PROLOG (PROgramming in LOGic) to encode the ontology and to query it, searching for logical inconsistencies in the theory. Once a domain ontology structure is developed and validated, segments can be incorporated as components of Web pages (as demonstrated in the previous section).

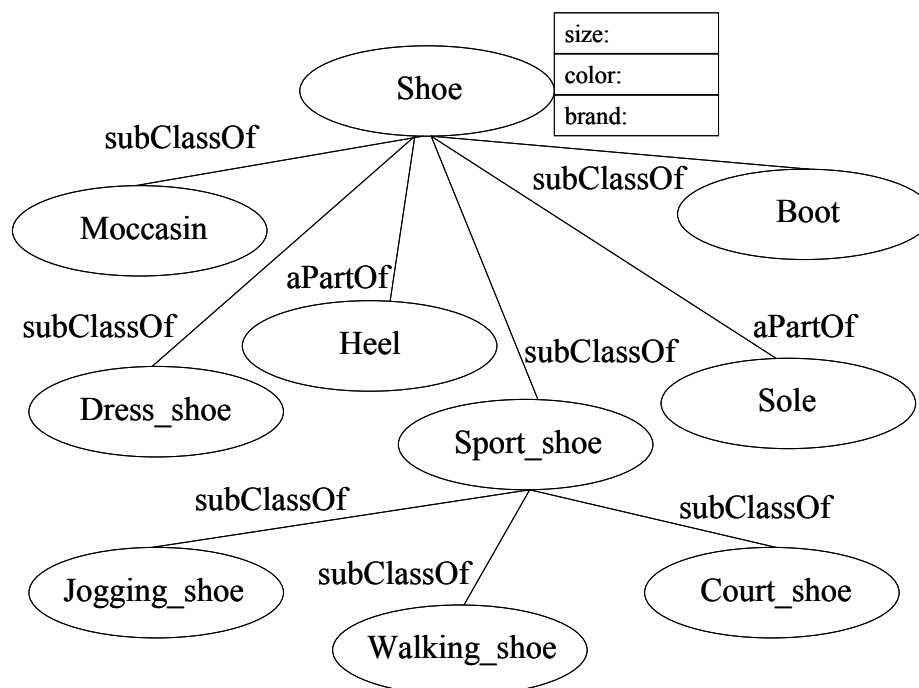


Figure 2.

A software agent can understand these segments and share that understanding with the human who gave the software agent a goal to achieve. Assuming the existence of several well-designed domain ontologies, the construction of a higher-level ontology is a matter of defining the "gates" between the different domain ontologies. For example, the "wears" relationship may be the gate through which a software agent might use both the Human ontology and the Shoe ontology.

4. UNDERGRADUATE RESEARCH

Since the tools and techniques to be used in the Semantic Web are an emerging technology, having an undergraduate computer course focus on how to use them is not appropriate at this time. However, we have found it very beneficial to use ontology development for the Semantic Web to promote undergraduate research. Unlike using research projects in undergraduate degree programs (Hollocks 2001), such as a capstone course or senior project, our approach is strictly voluntary (on the part of the faculty member and the undergraduate) and does not count for any student course work. Nonetheless, the value to the students involved is the same as Hollock expressed: (1) Understanding the concept of research, (2) Development of problem definition and goal skills, (3) Wider and more critical perspectives on literature sources, (4) Development of reasoning and inquiry skills, (5) Development of presentation and writing skills, (6) Development of self-management skills, and (7) Development of an in-depth understanding of a particular domain, which may support career opportunities after graduation.

Getting students interested in doing our type of undergraduate research is not an easy matter. First there is the need to motivate the research work. Students must be given a reason for going beyond course work, a well-defined and structured environment, into something new, subject to change, and needing structure (i.e., research). Given the vision of the Semantic Web and the knowledge that pieces of it (e.g., XML and the current Web itself) are in place and being used for e-Commerce, we can encourage students to develop small domain ontologies that may be embedded into the Semantic Web after

OWL becomes an established means for doing so. Second there is the problem of accessibility, that is, giving the undergraduate sufficient background (often not found in a required course) so that the research problem is understood and the research goal is clear. Again, if the focus is on developing an ontology then the problem of accessibility really becomes one of what domain the student is interested in working, a possible future career field. This has not proved to be a real problem; students have a plethora of domains in which they are interested, and it is more a problem of deciding to work on one small domain where research progress can be demonstrated in a reasonable amount of time, say, a year or two.

Thompson (2000), an undergraduate researcher, worked in the domain of high school algebra word problems. Her developed ontology was implemented in PROLOG and used by CASPOR (Computer Algebra Story Problem ORiginator), a PROLOG program that randomly generated specific types of algebra word problems (e.g., mixture) for students to solve. CASPOR used the ontology to represent the knowledge about the context of the word problems. For example, in distance-rate-time problems, a car object must be associated in context with an object such as a road. This knowledge context would prevent the generation of a problem that begins, "Two cars fly to Chicago at the same time." Instead, the ontology insures that the problem begins, "Two cars take the road to Chicago at the same time."

Another undergraduate researcher, Gilds (2002) worked in a more difficult domain, that of world religions. The various violent acts being perpetrated globally in the name of a religion motivated the development of this ontology. The student had a deep interest in her own religion and wanted to understand how other religions, that have common origins, could be in such conflict with one another. Religions are a psychosocial factor. Religions form a psychological context in which some people function violently, for example, the Kamikaze functioned in the context of Bushido, and the Hamas suicide bombers function in the context of Islam. Bushido and Islam are not as well related as say Islam and Judaism (both trace historical paths from Abraham). Religions also form a social structure within society; there are re-

ligious leaders (rabbi, imam, bishop, etc.), buildings (temple, mosque, church, etc.), and other trappings. All of these concepts and relationships between concepts have to be incorporated into a well-structured ontology to be understood by a software agent. On today's WWW there is no dearth of Web pages about religion and by various economic measures religion is big business. It is commendable that an undergraduate would even attempt to work in such a domain. However, after a year and a half of work, the PROLOG implementation of her

ontology has demonstrated that the ontology needs more work before a software agent can use it.

We conclude this section with a PROLOG implementation of the Shoe ontology developed in the previous section. Since one of the drivers for OWL is frame-based systems, we have our undergraduate researchers take a frame-based knowledge representation for the ontology. Specifically, we have them use the knowledge representation expressed in Table 2.

```
name( Concept ).
relation( Concept, Relationship, Another_Concept ).
slot( Concept, Attribute )
value( Concept, Attribute, What_it_is ).
```

Table 2

```
name(shoe).
name(moccasin).
name(dress_shoe).
name(sport_shoe).
name(boot).
name(jogging_shoe).
name(walking_shoe).
name(court_shoe).
name(heel).
name(sole).
relation(moccasin, subClassOf, shoe).
relation(dress_shoe, subClassOf, shoe).
relation(sport_shoe, subClassOf, shoe).
relation(boot, subClassOf, shoe).
relation(jogging_shoe, subClassOf, sport_shoe).
relation(walking_shoe, subClassOf, sport_shoe).
relation(court_shoe, subClassOf, sport_shoe).
relation(heel, aPartOf, shoe).
relation(sole, aPartOf, shoe).
slot(shoe, size).
slot(shoe, color).
slot(shoe, brand).

?- relation(What, aPartOf, Shoe).
What = heel;
What = sole;
no more solutions

?- has_slot(court_shoe, What).
What = size;
What = color;
What = brand;
no more solutions
```

What are the parts of a shoe?

What are the attributes of a court shoe?

NOTE: These attributes are inherited from the attributes of a shoe.

Table 3

The PROLOG implementation of the Shoe ontology is given in Table 3. Having a computer executable ontology allows undergraduate researchers to verify and validate their logical theory using simple queries. Examples of such queries and the program's response appear at the bottom of Table 3. The last query in Table 3 is generated by an

ontology axiom that is used to implement inheritance. The students consider such a PROLOG rule part of the ontology "engine" since it can be used with any ontology having the knowledge representation structure given in Table 2. For completeness, the PROLOG implementation for the inheritance axiom is presented in Table 4.

```
has_slot( Concept, Attribute ) :- slot( Concept, Attribute ), !.
has_slot( Concept, Attribute ) :-
    relation( Concept, subclassOf, SuperConcept),
    has_slot( SuperConcept, Attribute ).
```

Table 4

There is little doubt that undergraduates can do the work described in this section. The most challenging part of the student's research is getting the in-depth understanding of the domain and, with that understanding, a more critical perspective on the information sources used in the domain. Developing the ontology and implementing it in a descriptive first order logic programming language is not difficult. Querying the ontology looking for inconsistencies can be difficult and can show that the proposed ontology is not ready to have a software agent use it. However, all of this work, whether completely successful or not, is preparing, quite naturally, for the Semantic Web.

5. CONCLUSION

The Semantic Web is a technology that Tim Berners-Lee, the father of the WWW, is hoping to have in place by 2005 (Port 2002). It would enhance the current WWW so that software agents could understand concepts and relationships posted on the Semantic Web pages. Such machine understanding would improve e-Commerce. Computer information systems undergraduates cannot ignore such an emerging technology, yet the technology is not stable enough at this time to require course work. On the other hand, ontology development, a key component of this emerging technology, is stable enough to have students conduct undergraduate research work.

This paper first presented some of the underlying infrastructure being developed for the Semantic Web by W3C. An example was given how XML, RDF, RDFS, and DAML + OIL might be used to embed a logical theory

into a Semantic Web page. The paper transitioned to a focus on the logical theory that would give an explicit, partial account of a conceptualization (i.e., an ontology). An example was presented showing how an ontology might be developed in a specific domain. The paper suggested a type of undergraduate research as a vehicle for developing domain ontologies, implementing each in a first order descriptive programming language, and validating them via queries. Published examples of such undergraduate research were cited, and an implementation of the previous example ontology given. The thesis of the paper is that undergraduate research in ontology development is a natural way to prepare students for working on the future Semantic Web.

6. REFERENCES

- Berners-Lee, T. and M. Fischetti, 1999, *Weaving the Web: The Original Design and Ultimate Destiny of the World Wide Web by its Inventor*. Harper, San Francisco, CA.
- Berners-Lee, T., J. Hendler and O. Lassila, 2001, "The Semantic Web." *Scientific American*, 284, 5, 34-43.
- Boggs, R., 2002, "The X-Factor: Implications for Internet Programming Today and Tomorrow." In D. Colton, M. Payne, N. Bhatnagar, and C. Wortschek (Eds.) *The Proceedings of ISECON 2002*, v 19 (San Antonio, TX) 415b.
- Bowman, M., A. Lopez and G. Tecuci, 2001, "Ontology Development for Military Applications." *Proceedings of the Thirty-*

- ninth Annual ACM Southeast Conference (Athens, GA) 112-117.
- Chandrasekaran, B., J. Josephson and V. Benjamins, 1999, "What Are Ontologies, and Why Do We Need Them?" *IEEE Intelligent Systems*, 14, 1, 20-26.
- Cherry, S. 2002, "Weaving a Web of Ideas." *IEEE Spectrum*, 39, 9, 65-69.
- Davies, J., D. Fensel and F. Van Harmelen, 2003, *Towards the Semantic Web: Ontology-driven Knowledge Management*. John Wiley & Sons, LTD, West Sussex, England.
- Fensel, D., 2001, *Ontologies: A Silver Bullet for Knowledge Management and Electronic Commerce*. Springer-Verlag, Berlin, Germany.
- Gilds, B., 2002, "Knowledge Engineering: Ontology of World Religions." *The Journal of Computing Sciences in Colleges*, 18, 1, 304-308.
- Gomez-Perez, A. and O. Corcho, 2002, "Ontology Languages for the Semantic Web." *IEEE Intelligent Systems*, 17, 1, 54-60.
- Gruber, T., 1993, "A Translation Approach to Portable Ontology Specifications." *Knowledge Acquisition*, 5, 2, 199-220.
- Guarino, N. and P. Giaretta, 1995, "Ontologies and Knowledge Bases: Towards a Terminological Clarification." In *Towards Very Large Knowledge Bases: Knowledge Building and Knowledge Sharing*, N. Mars (Ed.), 25-32, ISO Press, Amsterdam, The Netherlands.
- Hollocks, B., 2001, "The Value of Research Projects in Undergraduate Information Systems Degrees." In D. Colton, S. Feather, M. Payne, and W. Tastle (Eds.) *The Proceedings of ISECON 2001*, v 18 (Cincinnati, OH) 14c.
- Lenat, D., 1995, "CYC: A Large-scale Investment in Knowledge Infrastructure." *Communications of the ACM*, 38, 11, 33-38.
- Lopez, A., 1993, "In Search of Meta-Knowledge." *Proceedings of the 1993 Goddard Conference on Space Applications of Artificial Intelligence (Greenbelt, MD)* 263-269.
- Lopez, A. and M. Saack, 1992, "Logic Programming and Metadata Specifications." *Telematics and Informatics*, 9, 3/4, 271-279.
- McGuinness, D., R. Fikes, J. Hendler and L. Stein, 2002, "DAML + OIL: An Ontology Language for the Semantic Web." *IEEE Intelligent Systems*, 17, 5, 72-80.
- Maedche, A., 2002, *Ontology Learning for the Semantic Web*. Kluwer Academic Publishers, Boston, MA.
- Noy, N. and C. Hafner, 1997, "The State of the Art in Ontology Design: A Survey and Comparative Review." *AI Magazine*, 18, 3, 53-74.
- Port, O. 2002, "The Next Web: Think the World Wide Web is a godsend? By 2005, Tim Berners-Lee aims to be replacing it with the Semantic Web, which will understand human language." *Business Week*, 3772, 96-102.
- Saacks-Giguette, M. and A. Lopez, 1993, "A Frame-based Design for TMS and CAMS Metadata for a Stennis Information Management System." *Journal of Systems and Software*, 20, 1, 87-92.
- Stojkovic, V. and W. Lupton, 2000, "Software Agents: A Contribution to Agent Specification." In D. Colton, J. Caouette, and B. Raggad (Eds.) *The Proceedings of ISECON 2000*, v 17 (Philadelphia, PA) 505.
- Thompson, T., 2000, "Ontology Development with CASPOR." *The Journal of Computing in Small Colleges*, 15, 3, 58-64.
- Valenti, A., T. Russ, R. MacGregor and W. Swartout, 1999, "Building and (Re)Using Ontology of Air Campaign Planning." *IEEE Intelligent Systems*, 14, 1, 27-36.
- van der Vet, P. and N. Mars, 1998, "Bottom-up Construction of Ontologies." *IEEE*

Transactions on Knowledge and Data Engineering, 10, 4, 513-526.

Weinstein, P. and G. Alloway, 1997, "Seed Ontologies: Growing Digital Libraries as Distributed, Intelligent Systems." Proceedings of the Second ACM Digital Libraries Conference (Philadelphia, PA) 91-99.

Welty, C. and J. Jenkins, 1999, "Formal Ontology for Subject." Journal of Knowledge and Data Engineering, 31, 2, 155-182.