

Agents and the Semantic Web

Alexander Harvey

CSCE896 – Artificial Intelligence

December 2008

Agents and the Semantic Web

- Brief overview of Semantic Web
- Ontologies
- Concentrating on the “Web” aspect
- Modeling/Tools Example
- Web Services and Web Agents
- Conclusions

Overview

What is the Semantic Web?

- The Semantic Web provides a syntactical and semantic infrastructure by which software agents may easily access, understand, and manipulate a network of data.

End Goal

- Software agents intelligently running around the web and performing complex actions for their users.

End Goal

- *“They painted a future of intelligent software agents that would head out on the World Wide Web and automatically book flights and hotels for our trips, update our medical records and give us a single, customized answer to a particular question without our having to search for information or pore through results.” 2*

Technologies for a vision

- A common language for representing data that could be understood by all kinds of software agents.
- Sets of statements that translate information from disparate databases into common terms.
- Rules that allow software agents to reason about the information described in those terms.

Resource Description Format

- RDF is the most fundamental building block for defining information on the Web.
- Each piece of data and any link connecting two pieces of data is identified by a URI.
- The two pieces and any notation indicating how they are connected, are grouped together into what is called a triple.

Ontologies

What is an Ontology?

- As it applies to AI, it's essentially the specification of a concept. That is, defined terms and relationships between them, usually in some formal and preferably machine-readable manner.

- *“I envision a complex Web of semantics ruled by the same sort of anarchy that rules the rest of the Web. Instead of a few large, complex, consistent ontologies that great numbers of users share, I see a great number of small ontological components consisting largely of pointers to each other.” 1*

Ontology Languages

- Individuals or groups may want to define terms and data they frequently use, as well as the relationships among those items.
- Web Ontology Language (known as OWL) is one standard that can be used to define ontologies so that they are compatible with and can be understood by RDF.

Concentrating on the “Web” Aspect

- The real power of the Semantic Web comes from the web aspect. The growth and adaptation should mimic that of the Web itself.

- A small number of tool creators will need to know the details, but most users shouldn't need to know that Web semantics exist to create machine-readable web content.

- Again, as the group ontologies grow from use they will coalesce into a more concrete infrastructure of data that can be easily collected and processed.

Modeling/Tools Break

Web Services and Web Agents

- On top of this infrastructure, agent-based computing will become much more practical.

- By creating a resource in an ontological language, a capable agent would be able to use the hierarchy (and property restrictions) to find matches through class and subclass properties or other semantic links.

- Even better, by using a combination of Web pointers, markup, and ontology languages we can also include a machine-readable description of a service (as to how it runs) and consequences of using the service.

- Agents could then autonomously find many possible ways to meet the users needs and, with Machine Learning, personalize a result.

Service Logics

Service logics can be created to govern interactions such as

- Between agents and service invocations,
- What an agent may and may not do,
- Communication amongst agents,
- As well as complex planning of a set of services that together achieve's a user's goal

Inference Engines

- These software programs examine different ontologies to find new relations among terms and data in them.
- Finding different relations among different sources is an important step toward revealing the “meaning” of information.

- Pic (Figure 4, pg 34 (1))

Real-World Applications

Drug Discovery

- A research team at Cincinnati Children's Hospital Medical Center is using Semantic Web technologies to integrate data from different medical databases in incompatible formats to develop tools to personalize and predict which drugs, and what doses, will work for a given individual (2)

SAPPHIRE

- Developed at the University of Texas Health Science Center at Houston, SAPPHIRE integrates a wide range of data from local health care providers, hospitals, environmental protection agencies and scientific literature to allow health officials to assess the information through different lenses, such as tracking the spread of influenza or the treatment of HIV cases.

DBpedia

- DBpedia is an effort to smartly link information within Wikipedia's nine million articles. This project will allow users to perform detailed searches of Wikipedia's content that are impossible today, such as, "Find me all the films nominated for a Best Picture Academy Award before 1990 that ran longer than three hours."

Conclusions

Need

- As the number of services grow and the specificity of our needs increases, the ability of current search engines to find the most appropriate services is strained to the limit.
- Semantic Web technologies help to remedy this by providing the infrastructure by which agents can find meaning in data and customize for our needs.

Advantages

- Disparate databases developed in isolation can be merged by referring to and relating between resources in any of them.
- The success of the Web in general (and information archiving sites like Wikipedia in particular) has shown that there is sufficient incentive to publish quality data to make the overall Web a useful and even essential structure.

Problems

- The openness has serious ramifications in the Semantic Web, which go beyond, for instance, considerations that were important for technologies like expert systems. **(3)**
- Inferencing must be done conservatively according to the Open World Assumption, because at any time new information could become available that could undermine conclusions that have already been made. **(3)**

References

1. Hendler, James. “Agents and the Semantic Web.” IEEE Intelligent Systems 16.2 (2001): 30-37.
2. Feigenbaum, Lee, et al. “The Semantic Web in Action.” Scientific American Dec. 2007: 90-97.
3. Allemang, Dean, and Jim Hendler. Semantic Web for the Working Ontologist. Burlington: Morgan Kaufmann Publishers, 2008.