Heather Dewey-Hagborg 10/25/05
Networked Expression

The Semantic Web and Intelligence

Structure and Possibilities of the Semantic Web

The Semantic web is essentially an internet of meaning. It is a way of creating a global network of content so that it is easily processable by machines. The problem with the web currently is that although a massive amount of information is available there is no global system for publishing the meaning of this information in a language computers understand.

The semantic web would flourish through the construction and implementation of autonomous agents; intelligent network parsers designed to make information search simple, intuitive, and fast. Instead of looking for the information we need based on syntactical content HTML guesswork, we could search for the meaning directly. Instead of going to a page and traversing a network of links to find what we need we could simply ask for it. The semantic web is intended to complement humans in areas in which they do not perform well, such as processing large volumes of information quickly or analyzing large texts for certain pieces of information.

The semantic web is based on established technologies such as eXtensible Markup Language (XML), Resource Description Framework (RDF), ontologies and intelligent agents. XML is the successor to HTML. It is an open technology that allows specification of the meaning behind the content. XML schema explain structure by defining shared vocabularies and providing the links to associate semantics with them. RDF provides meaning to the structure of XML documents. It expresses the meaning and relationships between different web pages and concepts by explicitly coding the structure of things, properties and values. For example, Heather Dewey-Hagborg (thing) is the writer of (property) computer programs (value). Subject, object and verb (or thing, property and value) are encoded in the document through a uniform resource identifier (URI) which ensures that the words on the document are linked to a unique and definition.

Ontologies go one step further, providing the relationships between the terms themselves. (i.e. term A on my web page is expressing the same concept as term B on your web page). An ontology is a file that formally defines relations among terms, for example, a set of classifications and inference rules. By providing indexes of meaning, ontologies provide the ability to search for specific concepts rather than keywords.

Though the logical language for making inferences based on RDF has not been developed and there is presently no standardized approach for associating RDF compatible metadata with HTML, theoretically the semantic web combined with other networking technologies will create a knowledge base for an artificial intelligence system distributed throughout the web, a system increasingly referred to as the "Global Brain". This point was highlighted well in the discussion of ontologies in The Scientific American article on the semantic web.(http://www.starlab.vub.ac.be/teaching/berners-lee.pdf) As they point out, through the exchange of ontologies agents could bootstrap their knowledge and reasoning abilities. To me this is the most exciting part of the whole semantic web idea; the possibility of creating intelligent agents which could adapt, learn

and evolve over time based on a decentralized knowledge network. A global super-organism which could access humanity's collective knowledge, analyze across linguistic barriers, and discover ideas and relations we never could have found on our own. The question this raises for me is whether the most obvious method of doing this, creating a traditional expert system style AI, is the best or most interesting solution.

Knowledge Based AI

Artificial Intelligence programs that achieve expert-level competence in solving problems by utilizing a body of knowledge about specific tasks are called knowledge based or expert systems. The fundamental parts of an expert system are the knowledge representation system, the knowledge base, and the problem solving model. The knowledge representation system determines how information is represented, the knowledge base is the physical body or database of that knowledge, and the problem solving model determines the inference rules and heuristics that will be used to extract the desired information.

The fundamental hypothesis of knowledge based AI is that intelligent behavior can be achieved through symbol manipulation by a computer. That logic and inference are enough to create intelligence. For pattern recognition tasks and data retrieval knowledge based systems generally do a good job, but is this true intelligence?

Problems

One of the main problems with expert systems is that in order to be effective they require huge amounts of knowledge, and since all of the knowledge of the system must be explicitly programmed, in order to get that knowledge someone has to do the data entry. Creating a specialized knowledge based system is feasible, but a generalized intelligence is a daunting and presently unrealized task. The largest experiment yet in symbolic AI is CYC. Started in 1984 by Douglas Lenat, the goal was to create the mother of all knowledge bases for common sense information, providing a foundation for future generations of expert systems. The CYC project involves hand-coding many millions of assertions. By the end of the first six years, over one million assertions had been entered manually into the knowledge base. It is estimated that it will require 2 centuries of work to increase this figure to the 100 million assertions that Lenat believes are necessary for CYC to begin learning for itself. That is an absurd amount of time.

Another problem I personally have with knowledge based systems is their basis in the computational theory of mind. The basic premise of the computational theory of mind is that intentional states such as beliefs and desires are relations between a thinker and symbolic representations of the state itself. In other words, we think in symbolic representations. This means that essentially all of our cognitive abilities are formalizable and computable, a statement I find questionable. Is all of our knowledge and behavior really reducible to a series of production rules? Or is there something more complex at work? Something emergent which cannot be described or necessarily predicted by examining the low level components?

Finally, formal logic-based AI simply hasn't worked very well historically. It is extremely fussy needing unambiguous systems of input data in order to function. This doesn't mesh very well with the constantly changing world we inhabit, and these systems are notoriously poor at updating. They don't learn from experience or their environment, and are generally fragile. Solution discovery grows exponentially long as the amount of

input data increases, meaning the more complete a system is and the more thorough its heuristic is for finding a solution, the slower it will be to find that solution. In summary, knowledge based systems are slow, fragile, stagnant, and don't exhibit generalized intelligence.

The Knowledge Base and the Web

Using a distributed and dynamic knowledge base like the internet could potentially solve at least the problems of getting the input data into the system and keeping it up to date. If even half of the websites currently online provided an ontology that would be over a billion documents of relations. Based on sheer numbers alone certainly something interesting would come of it. Since humans maintain their own websites the update process would be automatic, built in to the structure of the data itself. Human updates could combine with agent updates based on other sources including physical sensors, web traffic, and heavily traversed connections to keep the whole system dynamic and current.

Potential problems with this vision include creating a set of production rules which are robust and flexible enough to gracefully handle the changing information, and developing a heuristic method for traversing the massive network of information. Another issue is evaluation of information value, in other words, avoiding a tyranny of the rich and popular. If agents are adapting pages and links, identifying connections between related concepts, enhancing popular links and gradually fading those which are rarely used, how can we avoid a silencing of the minority opinion? Finally there is also the fear that the global brain could act as an autonomous digital dictator, making individuals secondary to the whole, consuming our knowledge and enslaving us as insignificant analog I/O.

Alternate Possibilities

Ultimately the intelligence of the semantic web will lie in its agents, and I am sure there will be a wealth of options to chose from. Perhaps the default agent, the built in MS-NBC-Fox-SBC-Yahoo-Walmart bot will be limited and biased, but I am convinced the open source community will develop agents to combat this trend representing a wide spectrum of interests. Over time the tools for development will grow easier to use and making your own bot from scratch will be as simple as making a blog.

But I am discussing two parallel lines of interest here, one being the need for intelligent effective agents who "serve their humans" well, the other being the creation of a truly autonomous artificial intelligence, independent of its contribution to human efficiency. Though I think the global brain knowledge base intelligence would be very smart and useful I don't think it would qualify as true autonomous intelligence. So what would? What techniques could we use to more fully explore the intelligence potential of the semantic web?

I believe through a combination of artificial neural networks, genetic algorithms, and agent based distributed processing we could create a system where each networked person has a hub of independent agents and the interaction of all these agents with each other, with humans, and with the surrounding environment creates an emergent, distributed, community-like intelligence. Akin to the idea of building a new city upon a giant field of grass and leaving the creation of paths between points up to the residents,

this new internet could be a wilderness ripe for discovery but also generally designed to suit your day to day traversal needs. From the main path you could always tell there were these connected tangential areas, they would never fade from your vision, but the traversal would be more of an adventure than a daily routine. No centralization would exist, rather the smarts would be embedded in the behavior of the system as a whole, a higher level characteristic, observable but irreducible.

Physically what would this look like? The idea plays out overtime, starting with small groups of agents and evolving into an autonomous virtual world. Imagine each person as a town, a small hub of activity between multiple agents customized for their interests. I might have 50 different agents specialized for my varied interests, each living independently but interacting with my other agents constantly. Each agent is simple, contains a neural network, and has a small set of life tasks or goals to concentrate on. Everyone's agents interact with each other and also interact with the external world through sensors embedded in our portable and wearable devices. They get instructions and queries from us, but also keep up with the world. They follow news relevant to their interests, they sense things like weather, traffic, and pollution in addition to sensing *us*, our interaction with others, our moods and our health, our physical and emotional well being. They learn from our behavior and life experience at an expedited pace. Over time agents discover other agents with similar interests and paths; they become friends and eventually have the option of reproducing. With the passage of time the entire agent world begins to evolve based on input from us and the environment. They grow increasingly autonomous, need less and less direct input from us, increasingly able to simply hang out *with* us. They become our friends and advisers as well as our agents. They form what from our current perspective is a virtual world, but from the future position will be seamlessly integrated with our world; no longer any distinction between what is physical and what is virtual.