# Web 3.0: The Dawn of Semantic Search

James Hendler, *Rensselaer Polytechnic Institute*

**Emerging Web 3.0 applications use semantic technologies to augment the underlying Web system's functionalities.**

In the past two January editions of this *Computer* column, I've had the pleasure of writing about the status of the Semantic Web, and particularly of its applied use in Web applications, increasingly coming to be known as Web 3.0. I'm happy to say development and deployment continue apace, and that for those of us who know where to look, we see a lot of progress. Of course, if we were really being successful, you wouldn't have to know what rocks to look under—it would be everywhere. Or would it?

## AN INFRASTRUCTURE TECHNOLOGY

One of the difficulties in explaining Web 3.0 is that, unlike the original Web browser or later Web 2.0 systems, Semantic Web technology tends to be an infrastructure technology. While Web companies are working to produce new and scalable tools, academic researchers are pushing the size and speed of Semantic Web back-end operations.

### Corporate development

Web developers are learning that they can build an application from a combination of traditional databases with RDF (Resource Description Framework) triple stores, the databases of the Semantic Web. Traditional databases provide scaling for the back-end dynamics that are well-developed and clear; semantic databases provide new functionality that requires Web linking, flexible representation, and external access APIs. Relational databases provide the computational beef, the triple stores the secret Web 3.0 sauce.

In fact, you've probably visited a website sometime in the past few weeks that was built this way, but as hardly anyone is using any kinds of "Web 3.0 inside" labels, it's not surprising that you didn't know.

The first generation of enterprise Web 3.0 systems uses behind-the-scenes "structural" semantics to extend their current capabilities—for example, taxonomies with simple properties that can be used to relate terms to each other or to integrate terminologies from multiple sites. This sort of "controlled vocabulary" has been around for a long time, but emerging technologies allow it to be more easily integrated with Web development (J. Hendler, "Web 3.0 Emerging," *Computer*, Jan. 2009, pp. 111-113).

In addition, new standards make it possible to find consulting and tool-development companies that can help provide the scalable back ends needed to make the systems succeed. One company I work with has a customer who has hired them to develop a "trillion triple" store that can keep up with its complex application's real-time needs.

### Academic research

While details about corporate use of the Semantic Web and the architecture to support it are still under wraps, the academic community is also looking more seriously at scalable reasoning and large-scale back-end applications.

Researchers are exploring the scaling of both triple-store capabilities and inference algorithms for Semantic Web languages. The best paper at the 2007 VLDB conference was on using a DBMS for Semantic Web data management (D.J. Abadi et al., "Scalable Semantic Web Data Management Using Vertical Partitioning," *Proc. 33rd Conf. Very Large Data Bases*, VLDB Endowment, 2007, pp. 411-422), and that work has led to a number of new techniques for accessing and optimizing Semantic Web data.

At last year's International Semantic Web Conference (*The Semantic Web—ISWC 2009*, LNCS 5823, Springer, 2009), researchers presented findings on using parallel architectures for performing reasoning over semantic Web data at scale, including the design of both a


**Figure 1. Real-time search engine TipTop combines language technologies with search to classify Twitter results into positive responses (green), negative responses (red), and other opinions (blue).**

MapReduce mechanism (J. Urbani et al., "Scalable Distributed Reasoning Using MapReduce," pp. 634-649) and a cluster-based technique (J. Weaver and J. Hendler, "Parallel Materialization of the Finite RDFS Closure for Hundreds of Millions of Triples," pp. 682-697) for computing the inference-based closures for more than a billon triples. Other work looked at more efficient reasoning in the space of OWL reasoning (J. Du et al., "A Decomposition-Based Approach to Optimizing Conjunctive Query Answering in OWL DL," pp. 146-162) at larger scales.

## SEMANTIC SEARCH

Web companies haven't been waiting around for research results, and they're starting to deploy specialized algorithms to meet their own needs. In fact, some new technologies are beginning to emerge from the Semantic Web infrastructure, and in 2010 we'll see more of these companies providing greater functionality to their users.

The most important area where we'll see these technologies on the Web is in the growing area of semantic search engines. These include systems that try to augment general searches as well as systems that are trying to literally change the search experience.

While the internal details of most of these systems are still proprietary, in general they appear to combine a pragmatic approach to natural-language processing with a lightweight semantics that lets them better collect and process information about specific areas. A complete survey is beyond the mandate of this column, but a few applications will suffice to highlight some of the differences between these and traditional systems.

### More informative results

One semantic search capability is the attempt to provide more informative results than are typically returned by a regular search engine. Rather than simply identifying a useful page, these systems try to pull the information from those pages that might be what a user is looking for, and to make this immediately apparent.

For example, a search for "James Hendler" with Sensebot (www.sensebot.net), one of the newer search engines, will return a set of results such as

James Hendler (born April 2, 1957) is an artificial intelligence researcher at Rensselaer Polytechnic Institute, USA, and one of the originators of the Semantic Web.

[*SOURCE:* ***James Hendler*** *facts - Freebase (www.freebase.com/view/en/james_hendler)*]

Sensebot uses language technologies to identify specific assertions about the object being searched for—my name in this case—and to provide a source for those assertions: It found this sentence in the Freebase open-database system.

## Further search suggestions

A second capability offered by semantic search is to try to help a user identify further searches that may be more useful, and which can identify related searches to help users hone in on what they're looking for.

Probably the best known example of this capability is offered by Microsoft's Bing search engine. For example, a Bing search for "IEEE Computer" returns, along with regular search results, results for a list of related queries such as "IEEE PCS" or "IEEE Computer Security Conference" that might lead a user to more detailed information.

These expansions vary in quality based on how much data Bing has on the particular thing being searched for and sometimes can be quite impressive. For example, a search for the actress "Gates McFadden" returns a sidebar of related results for Brent Spiner, Lavar Burton, Will Wheaton, and some of her other *Star Trek: The Next Generation* costars.

## "Affective" Web content

Another illustration of things to come in semantic search is the combination of language technologies with search to discover "affective" aspects of Web content, especially in the blogosphere or in Twitter feeds.

An example of such technology is the real-time search engine TipTop (http://feeltiptop.com). Figure 1 shows the results of a TipTop search for "web 3.0," which sorts several recent Twitter messages that mention the topic into positive responses (green), negative responses (red), and other opinions (blue).

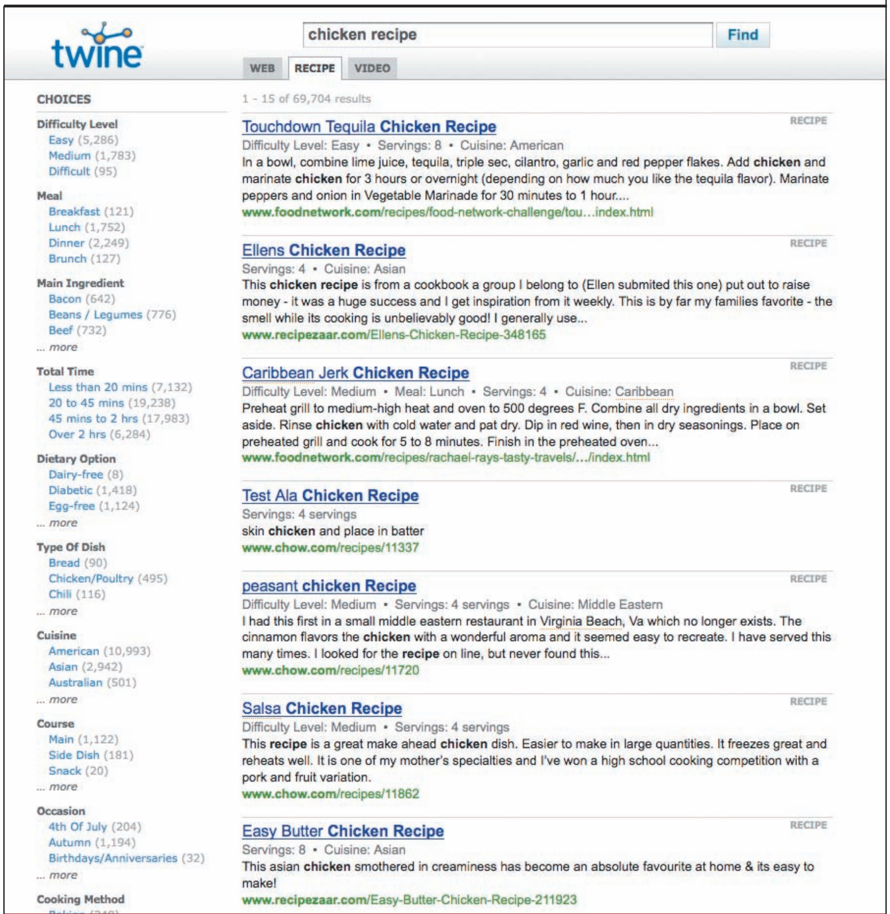


**Figure 2. Search engine T2 draws on domain knowledge to find and categorize search results.**

## Domain knowledge

An important use of semantics in search is to draw on domain knowledge in areas where searches are difficult.

At ISWC 09, Nova Spivack, CEO and founder of Radar Networks (developers of Twine), gave some examples of this capability in his company's forthcoming T2 semantic search engine. Figure 2 shows the results of a search for "chicken recipe." The engine uses semantic technologies not only to find recipes from sites but also to filter them by several categories including cooking time, dietary options, and cuisine.

Semantic search techniques that use domain knowledge clearly would change the search experience if widely deployed. However, it will take time and a combination of human and machine effort to cover the enormous diversity of Web domains. T2's solution to this problem is to provide the means for people to create these mappings using social, wiki-like mechanisms, thus extending the search engine's reach.

## Matching people and needs

Beyond simple keyword matching, another use of semantics and language technologies is to find matches between people and their needs.

A good example of this is Applied Informatics' TrialX (http://trialx.com), which won the 2009 Semantic Web Challenge (http://challenge.semanticweb.org). This application uses advanced medical ontologies to combine electronic health records with user-generated information to match people with potentially helpful clinical trials. Thus far, TrialX has successfully matched more than 3,000 participants to appropriate trials, helping both users and researchers in the pursuit of new medical treatments.

Another example of semantic matching is offered by Bintro (www.bintro.com), a classifieds website that matches users to whatever they are looking for—jobs, volunteer organizations to join, business partners, and so on—without them having to fill out a form or hope that someone notices them in a sea of static listings, or doing multiple keyword searches. Bintro uses a combination of semantic techniques to determine the meaning of a user's descriptions and then, through an asynchronous match process, finds other users with similar descriptions. Thus, a person looking for a position as a "childcare provider in the New York City area" could be matched with someone advertising for a "nanny in Manhattan."

Like most semantic search applications, Bintro takes advantage of Semantic Web standards to build independent ontologies and thereby rapidly move into new domain areas.

Emerging Web 3.0 applications are enhancing search engines in interesting new ways. Architecturally, they all have one thing in common—they use semantics to augment the underlying Web system's functionalities. When the semantics aren't applicable, or where they fail to add value, the underlying application looks like a traditional Web or Web 2.0 site. But where the semantics can be useful, the new functionality adds some exciting oomph. And on the Web, a little oomph can lead to a lot of money. ■

*James Hendler* is the Tetherless World Chair of Information Technology and Web Science at Rensselaer Polytechnic Institute, as well as an advisory or board member for numerous Semantic Web companies including Radar Networks and Bintro. Contact him at hendler@cs.rpi.edu.

**Editor: Simon S.Y. Shim, Dept. of Computer Engineering, San Jose State Univ., San Jose, CA; simon.shim@sjsu.edu**