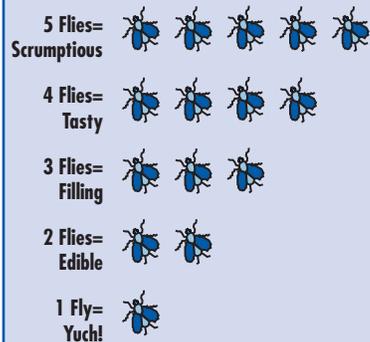


THE ARACHNOID TOURIST

Robert Filman • Lockheed Martin Advanced Technology Center • filman@ict.atc.lmco.com
Feniosky Peña-Mora • Massachusetts Institute of Technology • feniosky@mit.edu

ARACHNOID AGENT, AT YOUR SERVICE



One amusing thing about the Web is that it gives AI programs something “real” to work with. That is, it allows a purely software system to deal with a nonartificial world.

Traditional AI applications are either marked by artificially restricted problems or dominated by messy mechanical things like cameras and motors. It’s hard to make a robot that can roam the Earth effectively and economically—what with all the hardware to buy and maintain. And anyway, Mars rovers and the like aside, what can a robot do that a \$20-per-hour human can’t do just as well? Robotics is further complicated by irrational things like pedestrians—who wants to spend their research budget on liability insurance?

Cyberspace, on the other hand, is ideal for AI. There’s a lot of real stuff out there for a software agent to muddle through—far more than humans can handle—and it’s difficult to draw blood over the Internet.

So the Spider set out to look at agents on the Web. (We note that not every researcher working on agents thinks they are related to AI. After all, “agents” has become too hyped a term to have a very specific meaning. However, like so many computer anthropomorphisms, the term itself suggests intelligence.)

AGENT THEORY

Agents are popular. There are books on agents, conferences on agents, and columns on agents. Occasionally, a publication even gets the weird idea of devoting an issue to agents.

Intelligent Software Agents • www.cs.umbc.edu/agents/
Tim Finin
Visited 28 June 1997



So where to start? Probably the best place is Tim Finin’s “Intelligent Software Agents” page. Finin maintains the

agents mailing list, an obvious segue into an agents page, and is also the originator of KQML (www.cs.umbc.edu/kqml/), a language/protocol for information exchange.

At the site we found a “What’s New” section (books, conferences, and products dominate), a “Webletter” (more of the same), pages on the theory and technology of agents, directories of projects, agent applications, and agents in relation to other fields. (We are particularly fond of the agent humor.) In fact, we found everything our hearts desired—except a search engine to index the site.

Other good sites for agent indices include Botspot (www.botspot.com) and Agent Knowledge-Base Associates (www.akainc.com/tten.htm).

SHOE: Simple HTML Ontology Extensions •

www.cs.umd.edu/projects/plus/SHOE/

Sean Luke, Lee Spector, David Rager, and Jim Hendler
Visited 29 June 1997



The most common application of agents is search. But barring a clever natural language understanding system, searching systems are restricted to matching on syntactic keys—the presence or proximity of words—rather than any deeper meaning. The SHOE approach to this problem is an HTML extension that lets you annotate Web pages with semantic descriptors. One perpetual tension in AI has been between representation schemes that are computationally efficient but not particularly expressive (for example, name-value pairs) versus representations that are computationally inefficient (or even intractable) but more expressive (for example, full first-order predicate calculus).

SHOE picks the Prolog spot on this scale, allowing implications, subclasses, and universal quantifiers (Horn-clauses), but omitting negation and serious disjunction. This makes it computationally reasonable yet still able to communicate more than simple base-level facts. In the current spirit of ontologically driven AI, the authors are also working on various predefined representation schemata that can be inherited to particular pages.

The site includes the “Knowledge Annotator” (a Java applet) to help define the knowledge embedded in Web pages. It works relatively well, generating new HTML tags within the pages. In HTML, what the browser doesn’t understand it ignores; so ordinary browsers have no trouble viewing the extended pages. The authors also provide a path for loading SHOE data into the Parallel Representation of Knowledge and Associations (PARKA) representation/reasoning system for search and inference.

Unless they are widely adapted, things like SHOE remain interesting research curiosities. One advantage of the original WWW was that it incorporated existing services like Gopher and mail into the emerging HTML browsers. (Even so, the Tourist confesses surprise that the

Web caught on, no less at how quickly.) Working with HTML is a clear plus for SHOE. However, the Tourist wonders if using an existing (and more popular) representational language (like, say, a modified Prolog) might not result in a quicker adoption of the technology. Similarly, the existence and standardization of SHOE on enough reasonable, inheritable ontologies will be a major determinant of its success.

Aglets • www.trl.ibm.co.jp/aglets/

IBM Tokyo Research Laboratory

Visited 29 June 1997



One of Java's distinguishing features is its network orientation, including the ability to transmit code in a machine-independent way. This meshes well with the notion of mobile agents that transport themselves around a network. Aglets are IBM's Java foundation classes for such programs. They include primitives for communication and agent migration.

The site includes a working draft of several chapters for an emerging book, *Programming Mobile Agents in Java with the Java Aglet API* by Danny B. Lange and Mitsuru Oshima. But the Tourist's final impression of the site was rather like that of the four-year-married, still-virginal bride of the marketeer, who falls asleep every night to her husband's description of how wonderful it will be. There are several links to articles extolling the virtues of the impending mobile agents, a book detailing the API of these agents, and nary an example of an actual system doing anything that takes advantage of the agent mobility.

Aglet software is also available (free) from the site.

AGENT APPLICATIONS

The Spider moves along from theory to practice, and visits a few sites with agents in place.

CSTaR • www.ac.com/cstar/hsil/agents/

Chad E. Burkey, Theodore D. Anagnost, Benjamin E. Hosken, and Mark Jacobson

Visited 5 July 1997



This site reports on the agents work at Andersen Consulting's Center for Strategic Technology Research. Since Arthur Andersen is in the business-consulting business, it's not terribly surprising to find that its research centers on, well, business. The site presents two demos of agents: LifestyleFinder and BargainFinder.

The LifestyleFinder seeks the holy grail of marketing—the ability to infer from ancillary information what advertisements will best suit which customers—that is, to figure out that customers interested in ski masks are good targets for ads about skis and handguns. The nominal social good of this activity is that since the Web is supported by advertisements, you're going to get a fair volume of advertising thrown at you when surfing. Far better if the ads interest you.

This leaves us a little uncomfortable, since we believe that the collected information will more likely be used to benefit the information gatherer than the information provider. But being good tourists, we each answered the LifestyleFinder's questions and received, in return, a list of

15 recommended Web sites. Of these, Feniosky found three of possible interest. Bob was relieved to discover he was interested in exactly none. So even if "they" are secretly accumulating information, it's still noise to them. We leave it to others to debate to what extent the LifestyleFinder (which seems more like a questionnaire) is an agent.

From the name BargainFinder, the Spider had high hopes of pragmatic information, since the little hatchlings around here seem to break things faster than we can fix them. Alas, the BargainFinder finds bargains only for recorded music. Nevertheless game, the Spider picked an album at random and got back four quotes, three in the 13 dollars and change region, and one at \$10. Well, that's clearly useful. (Or at least it's useful right now with four choices. The interface will need a little work when the number of choices increases.) However, when we checked that \$10 quote, we found that it also came from a bargain-finder-like agent that presented six choices of where to buy it.

Left the Spider feeling a little like Dorothy being told to ignore the little man behind the curtain.

InfoSpiders • www.cs.ucsd.edu/~fil/agents

Filippo Menczer and Rik Belew

Visited 9 April 1997



InfoSpiders (aka ARACHNID: Adaptive Retrieval Agents Choosing Heuristic Neighborhoods for Information Discovery) is a Net search agent using endogenous fitness (an artificial-life-inspired algorithm) to search the Web. We're certainly charmed by the name. The idea itself (as described in the papers) is interesting, but we had great difficulty seeing it in action. The site has papers, the text of slide presentations, simulation demos (which we didn't download) and movies (which we tried to download but failed—we suspect bad address pointers).

The papers describe an implementation in which users provide keywords and the system breeds a set of agents that spring to life and begin crawling over the Web. Successful agents replicate and flourish; unsuccessful strains die out. We really wanted to see a demonstration. We wanted a page where we could play with the system and see a visual representation of the agents dying, reproducing, and surviving according to the different paths they take and the energy they expend or gather. What we got were some papers, slide presentations, and addresses of nonexistent movies.

Agent Info Space • www.argolink.com/agent

Visited 21 June 1997



This site provides a very interactive communication with a genie. The genie talks and supposedly listens. We weren't able to make it understand what we said, but were still intrigued by what we could see of the demonstration.

The example presented was the First Genie Internet Bank. The genie recommends which bills to pay, tells what the balance on the account is, and indicates when to request a credit increase. If the actual agent does everything the demo says it will do, banking will be a lot easier. As an additional fillip, the demo agent talks and gestures.

Be warned, though, that argolink is fickle. It seems quite Active X dependent, and on some visits we couldn't get it to work. ■