Cover Feature

The Internet, Intranets, and the AI Renaissance

Virtually cost-free publication on the Web has led to information overload. AI, with its roots in knowledge representation, is experiencing a renaissance as new tools emerge to make the Web more tractable.

Daniel E. O'Leary
University of Southern California

The Internet, fueled by the phenomenal popularity of the World Wide Web, has exhibited exponential growth over the past three years. But virtually instantaneous and cost-free publication, inherent in the WWW, leads to problems with information overload. In the case of the WWW, the ease of self-publication has helped generate an estimated 50-120 million pages on the Web, a figure that is growing every day.

Search engines help users navigate the millions of pages. However, even the best search engines cannot efficiently circumnavigate the entire Web. A naive search using AltaVista, for example, can result in 100,000-plus matches. New search engines can structure queries in a user-friendly fashion, but they require the user to learn and manage numerous interfaces. We may now have digital libraries on our desktops, but it still takes an immense amount of manual effort to use them.

AI will play a crucial role in making the WWW usable. AI has been around for 40 years, yet in some people's view it has not lived up to its (perhaps overblown) promise. Many domains create intractable problems for AI, and many expert systems are either too narrow or too brittle. But the Web is a perfect environment for AI, with its roots in problem solving and knowledge representation. An increasing number of Web-based AI applications—intelligent search engines and browsers, learning agents, and knowledge-sharing agents—have begun to emerge. And although these AI applications have grown increasingly attractive because of the Internet, they may hold even more promise for intranets.

INTRANETS AND THE INTERNET

Generic browsers, such as Netscape Navigator, can be used on intranets just as they are on the Internet. Market analysts are predicting that intranet growth, based on revenue, will leapfrog Internet growth by a factor of two between now and 2000.[1]

Although intranets employ classic HTML in a TCP/IP environment, they are more tractable than the Internet because their document format and content can be specified and controlled by the intranet's owner. In addition, intranets can employ protocols and database structures, such as Lotus Notes, that cannot be used in an Internet setting.

For all these reasons, intranets will provide fertile ground for an AI renaissance, perhaps even more fertile than the Internet. The market for intranet search and retrieval tools is expected to grow fivefold, from $53 million to $255 million over the next few years, according to Zona Research.[1] Already, Lotus Development and PointCast have joined forces to produce Domino.Broadcast, which speedily gathers business-critical data on the Web and disseminates it over an intranet, for the purposes of gaining a competitive advantage. For example, a Domino agent can search competitors' Web sites for price changes and then alert its own sales representatives, letting them adjust their prices accordingly.

AI SUPPORT FOR INDIVIDUAL ACTIVITIES

Some uses of AI on the Internet are aimed primarily at guiding the user, either offering direct assistance at the time of the activity (searching or browsing) or autonomous assistance in the background (independent agent search).

Search engines

In general, search engines help users find resources. In an Internet environment, typically there is time only for keyword searches. Internet search engines, therefore, employ minimal domain knowledge and adopt a very general user model (search-intensive). The inability to make assumptions about users, domains, and keywords limits a search engine's capacity to optimize the trade-off between missed and irrelevant pages.

In an intranet, on the other hand, search engines

0018-9162/97/$10.00 © 1997 IEEE

Although still in its infancy, research into agent technology has grown over the past decade, and agent technology is now beginning to appear in commercial products.

can anticipate user needs and effectively delimit connotations of keywords with multiple meanings. In addition, in an intranet environment the cost of false positive matches or missed documents is higher. Thus several corporations, including Eastman Kodak, are experimenting with building more knowledge and intelligence into their intranet search engines.

Arthur Andersen's FSA (Financial Statement Analyzer) and Eloise (English Language Oriented System for Edgar) systems search the US Security and Exchange Commission's Edgar financial database.[2] Both systems embed financial knowledge about particular concepts, and both model natural language understanding using *prediction* and *substantiation* in a manner similar to Gerald Dejong's work in Frump (Fast Reading Understanding and Memory Program).[3] In Frump, a *predictor* function makes predictions about what is likely to happen next, and a *substantiator* function verifies those predictions. For example, a search for "accounts receivable" does not provide all the different names by which accounts receivable might be listed in the database. But we can build additional knowledge into the predictor about other names that also signify accounts receivable, such as "trade accounts receivable." The substantiator can either substantiate a prediction based on content or find evidence for it based on concept inferencing. In a search for "accounts receivable," for example, prediction tests might compare the relative size of "accounts receivable" to the numerical quantities of other descriptors used in the prediction. In response to the question, "How many of this year's proxy statements contain changes to bylaws to create a new class of stock?" Eloise was able to correctly identify 85 percent of the cases contained in a test database.[2]

A similar set of issues underlies FAQFinder, from the AI Laboratory at the University of Chicago.[4] FAQFinder is an automated question-answering system that will take a query from a user and try to find the FAQ (Frequently Asked Questions) file that will most likely provide an answer. It uses a five-phased approach:

1. A statistically oriented information retrieval approach, called *Smart*,[5] finds FAQs that are possible matches.
2. FAQFinder parses the query into a syntax tree of simple noun and verb phrases, to obtain a representation used to support content matching.
3. FAQFinder's question recognizers operate on the parse tree to identify the appropriate category for the question. For example, questions of the type "What is the difference between..." are categorized as *q-comparison questions*.
4. FAQFinder performs semantic concept-matching using the *WordNet* network of lexical semantics and a thesaurus to select possible matches between the query and target questions.
5. FAQFinder presents matches to the user. If there is no close match, FAQFinder uses some heuristic question-matching strategies.

FAQFinder faces an extremely difficult and dynamic environment on the Internet, because FAQs can change substantially over time and there can be multiple FAQs for roughly the same domains. Intranet FAQ finders face a less complicated environment

Selected Web Sites

General-purpose intelligent agents

Autonomy Software—www.agentware.com
Browser Buddy—www.softbots.com

Special-purpose information-finding agents

Bargain Finder—http://bf.cstar.ac.com/bf
Firefly—http://www.agents-inc.com/agents/AgentsInc.html
NewsWeeder II—http://www.empirical.com
Webdoggie—http://webhound.www.media.mit.edu/projects/webhound

Intelligent browser assistance and support

Letizia—http://lcs.www.media.mit.edu/people/lieber/Lieberary/Letizia/Letizia-Intro.html
WebWatcher—http://www.cs.cmu.edu/afs/cs.cmu.edu/project/theo-6/web-agent/www/project-home.html

Intelligent browsing of Lotus Notes and large databases

ContactFinder—http://www.ac.com/cstar/hsil/agents/
FAQFinder—http://infolab.cs.uchicago.edu/faqfinder
Globenet—http://www.watson.ibm.com:8080
Scatter/Gather—http://theory.lcs.mit.edu:80/~karger/

Virtual groups

ARPA—http://dc.isx.com/projects/kiosk/aboutsisto/Descriptions.html
Dedal—http://buchu.eit.com:8888/dedal.html; http://gummo.stanford.edu/html/GCDK/dedal/dedal-info.html
Knowledge Sharing—http://www-ksl.stanford.edu/knowledge-sharing/papers/index.html
National Interchange for Knowledge Exchange—http://www.eit.com/creations/papers/nike/
ProcessLink and NextLink—http://cdr.stanford.edu/ProcessLink; http://cdr.stanford.edu/NextLink/
Web Librarian—http://www.eit.com:80/research/nasa.sbir/web_librarian.html

Virtual organizations

Agile Infrastructure for Manufacturing Systems (AIMS)—http://aims.parl. com/About-AIMS.html; http://www.eit.com/creations/papers/DMC93/
Defense Manufacturing Conference—http://www.wl.wpa.wpafb.af.mil/mtx

because organizations can limit the overlap between FAQs (for example, Human Resources would provide the only FAQ on pensions).

Finally, search algorithms that have been considered inefficient for large search spaces sometimes prove surprisingly effective for a relatively small set of documents (say, 5,000). Douglass Cutting and colleagues discuss a system called Scatter/Gather, which clusters data into semantically similar groups and presents it in a table of contents form. The table of contents includes groupings, summarized by "topical terms" and "typical titles," which can be further subcategorized and analyzed. Scatter/Gather works with a linear-time algorithm that can organize 5,000 documents in less than one minute on a Sparc20 workstation. The tool is particularly suited to vaguely defined information requests.[6]

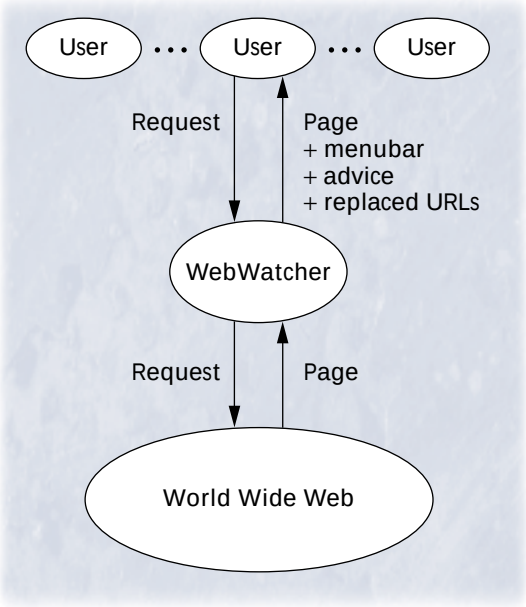

Figure 1. WebWatcher serves as an interface agent between the user and the Web.

Intelligent browsing

Although still in its infancy, research into agent technology has grown over the past decade, and agent technology is now beginning to appear in commercial products. Intelligent agents have been developed to directly facilitate browsing by learning the user's interests and providing a guided tour of the Internet. Two of the better known research prototypes include WebWatcher[7] and Letizia,[8] both from Carnegie Mellon University.

WebWatcher is a server-based interface agent that resides between the user and the Web as depicted in Figure 1. Any user running any browser can enter the system simply by typing a topic of interest in WebWatcher's FrontDoor page. WebWatcher accepts the request; replaces the current page with a modified page that embeds WebWatcher command menus and enables WebWatcher to follow the user as he browses; and presents to the user a highlighted list of recommended hyperlinks. Because WebWatcher is a server-based system, it can log data from thousands of users to "train" itself and refine its search knowledge. If a user signals that a particular search was successful, WebWatcher annotates each explored hyperlink with user interest keywords, adding to the knowledge base from previous tours.[7]

WebWatcher uses information-retrieval techniques based on the frequency of weighted terms and documents for all hyperlinks on a page, as well as user statistics associated with those links. WebWatcher can implement one of four learning methods:

- *popularity*—frequency of previous link traversal,
- *annotation*—relevance based on previous user interest,
- *match*—metric analysis of underlined anchor text, and
- *q-learning*—reinforcement learning that evaluates the value of downstream pages.

For each hyperlink it encounters, WebWatcher's algorithm ranks the associated list of interests in accordance with the user's stated interests and recommends a link if the ranked value crosses a certain threshold. In an experiment comparing WebWatcher's performance against human experts familiar with the pages used in the test suite, the system obtained an accuracy of almost 43 percent as opposed to the 47.5 percent achieved by the test subjects. The various learning methods achieve an accuracy that fluctuates according to whether the page is known or unknown, but which ranges between roughly 40 and 45 percent for all pages.

Henry Lieberman's Letizia is a client-side personal agent that collects information about the user's browsing habits and tries to anticipate additional items of interest.[8] Making inferences about user interests and using various heuristics, Letizia conducts a resource-limited search of the Web during idle time, looking for promising links to suggest when prompted.

Letizia uses a combination of information-retrieval and information-filtering strategies: Information filtering employs a passive user model, in which the system takes responsibility for removing less relevant material, while information-retrieval employs an active model, in which user queries are designed to select a set of relevant materials. Letizia's search analysis capabilities are presently limited to parsing a list of keywords and generating a preference ordering of interest among a set of links, although the Letizia and WebWatcher developers recognize the need for natural language capabilities that would capture syntactic and semantic information.[7,8] Letizia's search strategy uses a breadth-first search to compensate for the user's tendency to drill down through a hierarchy of Web documents, only to miss relevant neighboring documents at the top of the hierarchy.

Like search engines, intelligent browsers may play an even larger role in the corporate intranet. In their WebWatcher experiments, the authors expressed disappointment that users followed a recommended link

Figure 2. Autonomy Web Researcher interface. The user retrains the agent by telling it whether retrieved documents are relevant to a particular search.

only about 50 percent of the time. They explain this in part by users' short attention spans and by the diversity of user interests.[7] They speculate that their system would attain greater accuracy if it covered a narrower domain, which is precisely what intranets offer. In corporate settings, content can be specified and controlled, knowledge can be built in regarding relationships between different types of pages, and intelligent support for browsing can be used to point users in predetermined directions.

Intelligent agents

In contrast to intelligent search engines or intelligent browsers, intelligent agents perform tasks in the background while the user is performing other tasks. Thus, they need certain technical capabilities to perform these tasks independently. First, they must be able to match patterns and handle relatively complex logical comparisons. In addition, agents should have hierarchical (the relative relationship between concepts) and temporal intelligence. In particular, agents should be able to inherit rules from other structurally related agents and should be able to determine the current month in a rule like: "For the current month, which of the products have sales above budget?"

Furthermore, because personal agents are intended for individuals, they must be easily personalized and easy to build. Thus, problem- and opportunity-detection rules must be intuitive, and agents must interface with data in an intelligent manner. If data is not directly available but is derivable, the agent should recognize this fact and be able to generate such derivable data. Finally, when an agent comes back with an "alert," that information should be placed in context. Agents need to tell users why the alert was generated; where to go for further information without having to generate a complicated query; and, ideally, how to resolve the problem or capitalize on the opportunity.

General-purpose agents. Most agents depend on the user for definition and purpose. Browser Buddy is a rule-based agent that can be used to organize and access Web pages. For example, it can be used to access pages that would take a long time to interactively access or it can be set up to bring in information overnight, providing the user with a service analogous to the morning news.

Agents can also be learning agents. For example, Agentware recently released its Autonomy agent, which employs neural nets to search for information patterns rather than keywords. The user instructs the agent about a concept area using simple natural language explanations or descriptions. The agent's conceptual knowledge can subsequently be refined through a "retraining" exercise whereby the user recalls the agent and tells it which of the retrieved documents effectively describe the concept. Figure 2 shows the Web Researcher interface with Agent Fido off learning about AI. The heart of the architecture is Cambridge Neurodynamics' Dynamic Reasoning Engine, which combines probabilistic and neural network methods (including fuzzy logic) to perform data mining and analysis. Figure 3 illustrates the concept-clustering approach underlying this agent technology.

Special-purpose agents. A number of intelligent systems have been developed to find information or products on the WWW. Typically, these agents depend on the user's providing information either directly or indirectly through their choices so that the system can learn about their interests. Generally, these agents use rule-based approaches to search a large base of documents and distill that information. Four of the better known systems include

- Firefly uses *social information filtering* techniques[9] to help users find movies or music they are likely to enjoy on the basis of recommendations of others with "similar" tastes. User profiles are created using statistical correlations between an individual's preferences and data gathered from all users.
- CMU's Webdoggie attempts to mitigate the impact of information overload by recommending WWW documents based on user preferences (for example, whether users find them interesting or boring). Webdoggie uses rules based on explicit preferences or implicit choices.
- CMU's NewsWeeder II is a Mosaic-based news reader that learns users' interests and then uses machine learning to find new Web pages and news articles. (NewsWeeder is being commercially developed by Empirical Media Corporation.)
- NewsFinder[10] is an agent that routinely gathers online news based on user profiles. The system structures and executes the necessary queries for each individual news service, extracts articles that meet query requirements, determines which have not been seen before, and then matches the stories to the user's profile. It learns about a user's

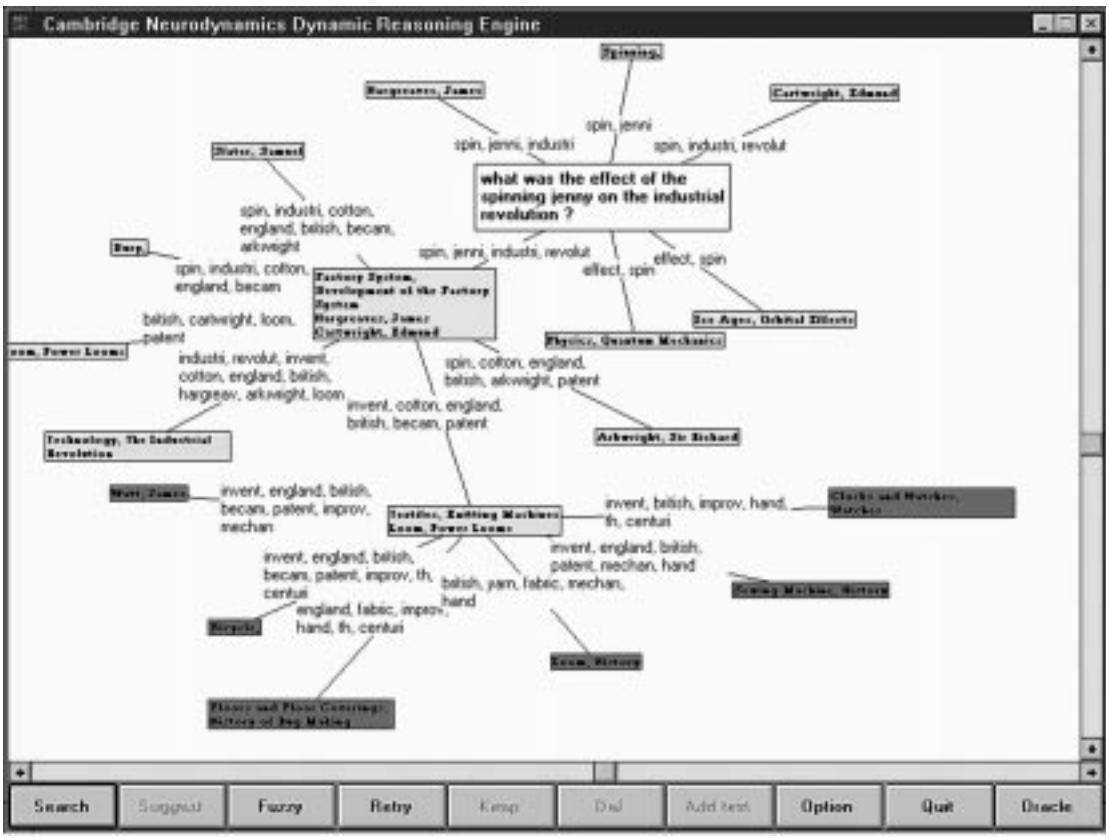

Figure 3. The Dynamic Reasoning Engine combines probabilistic and neural network methods (including fuzzy logic) to perform data mining and analysis. The system parses simple natural language queries and clusters the data in ways that might suggest other potentially relevant data to the user.

interests through a two-way dialogue, in which users tell NewsFinder which messages meet their interests and which do not.

- BargainFinder is designed to help the user find the "lowest" price for a CD.[10]

As intranets grow, agents can play an increasingly important role when they are freed from some of the problems that agents face on the Internet. For example, NewsFinder needs to know the different query requirements of each news source it investigates. As new sources become available, NewsFinder must learn new protocols. However, in an intranet environment, query formulation can be standardized so that systems do not need to be updated for new query formulation.

SUPPORTING DEPARTMENTAL ACTIVITY

Not all uses of AI and intelligent agents are aimed at supporting the decision making of individuals. Some uses are designed for broader groups.

Intelligent browsing of large databases

Although most browser research has focused on the WWW, there is interest in being able to browse corporate information repositories such as those developed in Lotus Notes. This will become even more important with the convergence of Internet-type architectures and more traditional information system architectures through network-centric products like Notes or Java applications.

This field has been of particular interest to major consulting firms, such as Andersen Consulting, who were among the first users of Notes. Despite the power of collaborative tools like Notes, they still require active user participation to fulfill their promise: A bulletin board intended for corporate problem solving is only effective if employees read it periodically and share their collective wisdom. Hence the attraction of automating the process. Bruce Krulwich's ContactFinder,[11] developed at Andersen Consulting's Center for Strategic Technology Research, is an agent-based system that monitors the company's internal bulletin boards that deal with technology, extracts questions, and refers questioners to employees perceived by ContactFinder to have relevant expertise.

ContactFinder focuses on information extraction rather than document understanding. It evaluates keyword, title, and subject fields for topic indicators and evaluates the name field for potential contacts. Moreover, it uses knowledge-free, or domain-independent, heuristics to find semantically significant phrases in both raw and rich text fields. The phrase "Does anyone…?" signals a question to be extracted. Short phrases in a different format, say italics, might signal something of import. So might numbered or bulleted lists, section headings, an acronym, or a compound noun phrase. The system is extremely conservative in its recommendations but has nonetheless read more than 5,000 messages and has referred 13 percent of those to "experts" with an accuracy of about 86 percent.[11]

Intelligent help-desks

IBM's Globenet is a knowledge-based information retrieval system currently deployed in a customer service support application that assists IBM staff in handling customer questions and problems.[12] Rule-based

Shared knowledge is at the core of organizational or group memory and is essential to the preservation of expertise or process knowledge.

agents control information retrieval from heterogeneous databases—in this case, network newsgroups. Agents are periodically sent out over the Internet to retrieve recent newsgroup postings pertaining to IBM products. Once the agent returns with that information, it uses rule-based knowledge to process it accordingly. An example rule might look like this: "If the Source is 'CompuServe' and the Newsgroup name is 'IBMPC/486/Software' and a contained Keyword is 'mobile' or 'PDA' then place in folder 'me/software/home-machines/laptop'." Rules can specify

- information sources to examine,
- information to present to the user (that is, which newsgroup and what key words),
- how to categorize the information (for example, which folder to put it in),
- how to prioritize the information (for example, how to order the information in the folder), and
- to whom to send the information.

Globenet reasons about both structured information (name of information sources or newsgroups, for example) and unstructured information in both headings and mail messages. In particular, Globenet can identify the existence of questions using a heuristic procedure and can do simple natural language processing by determining if the text includes particular keywords and phrases. Early tests with Globenet suggest that productivity was improved by more than 30 percent.[12]

Agent page monitoring

Organizations have found that paper is bulky and difficult to change. They are placing more material on their intranets, where employees can easily and rapidly find information about payroll statements, vacation policy, and the like.

There are other advantages to disseminating such information on the intranet. Firms like to make sure that employees read certain materials as part of the hiring process. Unfortunately, it is virtually impossible to determine if the appropriate material has been read. If material is accessible via the intranet, computer-based agents can be charged with watching pages and reporting what employees have been there and for how long.

Of course, there is a blurred line between normal monitoring activities and invasion of privacy. Because agents facilitate a broad-based ability to monitor employees, their use may imply certain value judgments and may require close scrutiny, particularly in the areas of privacy and security on intranets. For example, a female employee visiting pages relating to pregnancy leaves of absence might spur managerial inferences that she is pregnant. If such information is then used to limit promotion or bonuses, such monitoring would be an invasion of privacy and, ultimately, a violation of rights.

VIRTUAL ORGANIZATIONS

Virtual groups are distributed organizations and teams of people that meet and work together online. Group members rely on support systems to help gather, retrieve, and share relevant knowledge. Support systems enable the rapid assembly of groups or teams to solve particular problems, such as "virtual tiger teams" or proposal preparation teams consisting of individuals throughout the world.

Shared knowledge

Shared knowledge is at the core of organizational or group memory and is essential to the preservation of expertise or process knowledge. Accordingly, one of the primary issues in the case of virtual groups is using technology to share knowledge. ARPA's Intelligent Information Services project has moved to support virtual groups with a number of emerging technologies, including

- institutional memory tools that help organizations capture expertise, including process knowledge and access to expert consultants,
- tools to support multiuser/multiauthor hypermedia Web development so groups can build their own Web sites, and
- self-organizing knowledge repositories that adapt to community needs with use.

Stanford University's NextLink and ProcessLink projects have employed agent-based technology to enable distributed engineering groups to coordinate design decisions through peer-to-peer communication. Engineers make design decisions while agents do some of the bookkeeping by determining constraints and constraint violations. In particular, the systems use agents to coordinate constraint violations that occur with incremental revisions caused by constraint changes and preference changes by multiple engineers. Domain-specific agents from particular engineering domains coordinate through a constraint manager agent. Constraint manager agents monitor changes for such violations, using AI-based constraint satisfaction, and inform users when they occur.

NASA and Stanford's Generation and Conservation of Design Knowledge (GCDK) project focuses on methodologies and tools to capture design knowledge. An important component of GCDK was the Web Librarian project, which has as its goal the capture and retrieval (using Prolog rules) of large quantities of design information summarized in different forms, such as drawings, meeting notes, technical reports, and videotapes. Dedal, part of Web Librarian, provides mechanisms for indexing multimedia design information. In particular, Dedal uses model-based reasoning to enable queries based on topic, subject, medium (for example, video), and level of detail. It uses a domain-

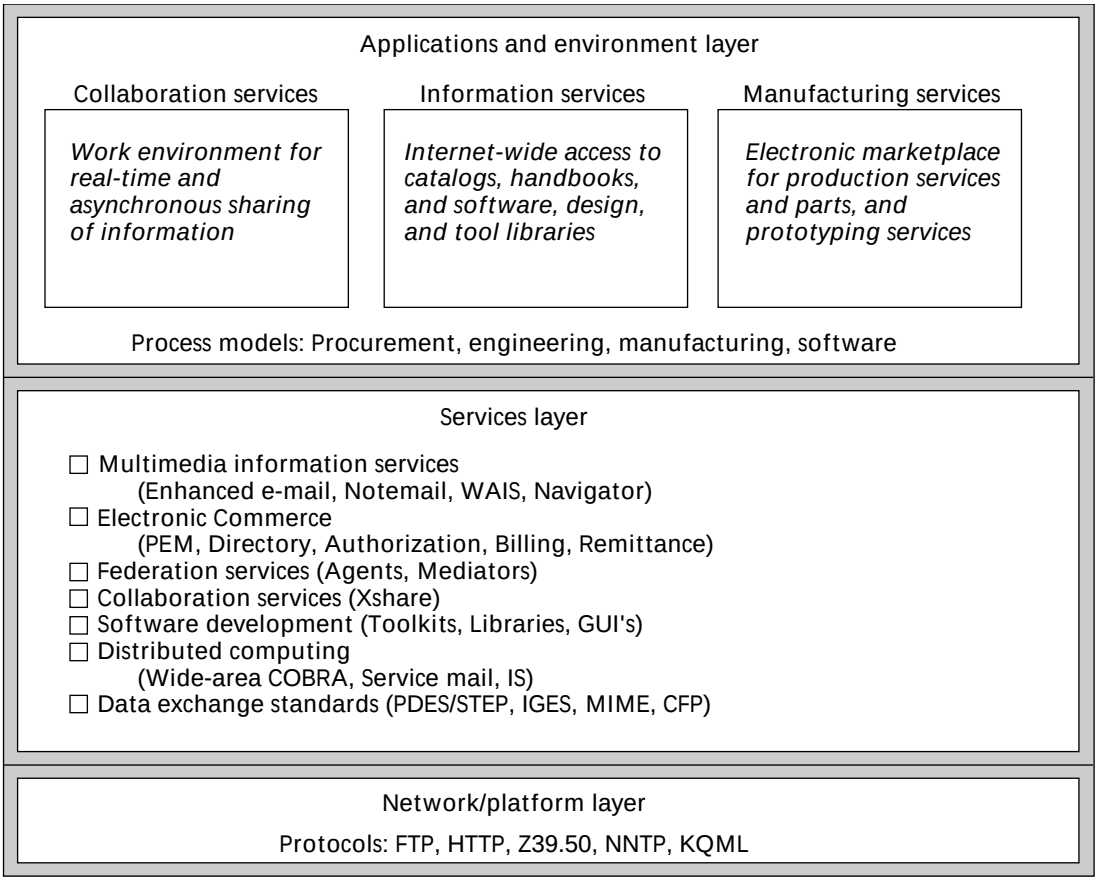

Figure 4. AIMSnet builds on existing networking and data-exchange standards to implement an open, scalable, distributed enterprise environment. AIMS-specific protocols for handling manufacturing and procurement data are layered on top of standard Internet communication protocols and are accessible via standard browsers.

based model and a set of retrieval heuristics to find responses to retrieval requests. If users find the retrieved information useful, the system acquires a new index.

Virtual companies: ARPA's AIMS project

Steven Goldman and colleagues define a virtual company as one "where complementary resources existing in a number of cooperating companies are left in place, but are integrated to support a particular product effort for as long as it is viable to do so…. Resources are selectively allocated to the virtual company if they are underused or if they can be profitably used there more than in the 'home' company."[13]

One of the more ambitious efforts to create a virtual organization is the ARPA/Air Force AIMS (Agile Infrastructure for Manufacturing Systems) project,[14] led by Lockheed Martin in conjunction with Texas Instruments, Rockwell, and other companies and institutions. The rationale behind the AIMS project is that a post-Cold War US can no longer afford a separate defense infrastructure and must instead develop a dual-purpose industrial base that can support defense needs and still compete globally. The three-year pilot project is intended to provide a model for agile manufacturing that can support dual-use technologies by providing high-quality, customized products quickly at the lowest cost. The Lockheed Launch Vehicle project will serve as the primary testbed, with manufacturers focusing initially on rapid prototyping and fast-turnaround small-lot production of precision machine parts for the LLV. As the AIMSnet matures over the next three to five years, it will add automated support for other fabrication and assembly operations and will expand beyond the testbed environment to other manufacturing initiatives across the country.

AIMSnet is a networked manufacturing and procurement infrastructure linking customers, suppliers, and other service brokers, using intelligent agents to carry out many business decisions. Because AIMSnet provides *certified manufacturing services*—meaning that business processes (contractual agreements and accounting practices, for example) and product data formats are also standardized—ontologies can be specified for autonomous agents, enabling substantial automation of manufacturing services. (Ontologies are specifications of discourse between multiple agents in the form of a shared vocabulary.)

AIMSnet builds on existing networking and data-exchange standards shown in Figure 4 to implement an open, scalable, distributed enterprise environment. AIMS-specific protocols for handling manufacturing

and procurement data are layered on top of standard Internet communication protocols, primarily secure HTTP, and are accessible via standard browsers. The system makes extensive use of interacting intelligent agents, whose roles include sending out requests for quotations (RFQ), determining if the unit has the spare resources to bid on the RFQ, putting together a bid on the RFQ, and choosing the bid that will be executed. Manufacturing agents can coordinate production schedules and balance loads among different vendors, and engineering database agents can notify each other of design changes that affect other members of the design team. The specific goals of the AIMS project are to reduce cycle time and costs by 25-30 percent.

Why do we say these Internet-based applications herald an AI renaissance? Artificial intelligence has come to play a crucial role in InfoAge retrieval strategies. Internet-based applications can exploit a wide range of AI developments. In this brief survey, we have seen examples of the following AI technologies:

- natural language processing (Concept-based Internet search);
- machine learning (WebWatcher)
- heuristic rules for establishing preference (Letizia)
- rule-based/heuristic natural language processing (ContactFinder, FAQFinder, Globenet); and
- neural networks (Autonomy).

This isn't AI for AI's sake—this renaissance is not one of stand-alone AI applications. Unlike first-generation AI applications, AI now can be embedded in heterogeneous networked computing environments and used for search, retrieval, and analysis of previously unimaginable quantities of data. Because the wealth of data makes direct human analysis impossible, AI-based support has become necessary to help users fully exploit that information.

Our increasingly competitive and technology-driven world has decreased the time available to us for decision making. To survive in this environment, we are increasingly turning to advanced computer technologies, such as intelligent agents, and delegating some of that decision making to these electronic surrogates. ❖

References

1. *Internet and Intranet 1996: Markets, Opportunities, and Trends*, Zona Research Inc., Redwood City, Calif., 1996.
2. C. Mui and W. McCarthy, *IEEE Expert*, Fall 1987, pp. 33-41.
3. G. Dejong, *Skimming Stories in REAL Time: An Experiment in Integrated Understanding*, doctoral dissertation, Dept. of Computer Science, Yale Univ., New Haven, Conn., 1979.
4. K. Hammond et al., "FAQFinder: A Case-Based Approach to Knowledge Navigation," *Proc. 11th Conf. Artificial Intelligence for Applications*, IEEE CS Press, Los Alamitos, Calif., 1994, pp. 80-86.
5. C. Buckley, "Implementation of the Smart Information Retrieval System," Tech. Report 85-686, Cornell Univ., Ithaca, NY, 1985.
6. D. Cutting et al., "Scatter/Gather: A Cluster-Based Approach to Browsing Large Document Collections," *Proc. Int'l SIG Information Retrieval*, ACM Press, New York, 1992.
7. T. Joachims, D. Freitag, and T. Mitchell, "WebWatcher: A Tour Guide for the World Wide Web," tech. report, Carnegie Mellon University, Pittsburgh, 1996; http://www.cs.cmu.edu/afs/cs.cmu.edu/project/theo-6/web-agent/www/project-home.html.
8. H. Lieberman, "Letizia: An Agent that Assists Web Browsing," 1995; http://lcs.www.media.mit.edu/people/lieber/Lieberary/Letizia/Letizia-Intro.html.
9. U. Shardanand and P. Maes, "Social Information Filtering: Algorithms for Automating 'Word of Mouth,'" *Proc. SIG Computer and Human Interaction*, ACM Press, New York, 1995; http://www.acm.org/sigchi/chi95/Electronic/ documnts/papers/us_bdy.htm.
10. B. Krulwich, "Information Integration Agents: BargainFinder and NewsFinder," Tech. Report WS-96-06, Am. Assoc. Artificial Intelligence, Palo Alto, Calif., 1996.
11. B. Krulwich and C. Burkey, "The ContactFinder Agent: Answering Bulletin Board Questions with Referrals," *Proc. Am. Assoc. Artificial Intelligence*, MIT Press, Cambridge, Mass., 1996, pp. 10-15.
12. B. Grosof and D. Foulger, "Globenet and Raise: Intelligent Agents for Networked Groups and Customer Service Support," 1995; http://www.watson.ibm.com: 8080.
13. S. Goldman, R. Nagel, and K. Preiss, *Agile Competitors and Virtual Organizations*, Van Nostrand Reinhold, New York, 1995.
14. H. Park, J. Tenenbaum, and R. Dove, "Agile Infrastructure for Manufacturing Systems (AIMS)," *Proc. Defense Manufacturing Conf.*, 1993; http://www.eit.com/creations/papers/DMC93/DMC93.html.

Daniel E. O'Leary is an associate professor at the University of Southern California. His research interests include the impact of intelligent agents on individuals, organizations, and commerce, and the integration of AI and reengineering efforts. O'Leary is the current editor, and former associate editor, of IEEE Expert. *He received a PhD from Case Western Reserve University.*

Contact O'Leary at 3660 Trousdale Parkway, University of Southern California, Los Angeles, CA 90089-1421 or e-mail oleary@rcf.usc.edu.