A Collaborative Fuzzy Expert System for the Web

Tod A. Sedbrook

University of Northern Colorado

Abstract

A convergence of Internet and fuzzy logic technologies provides an opportunity for experts and end users to collaborate in developing, refining, and testing knowledge-based systems. Internet technology removes geographical and time-based restraints, and fuzzy rule bases are easier to understand and maintain. This paper describes an architecture and a prototype for developing, delivering, and maintaining expert systems on the World Wide Web.

The system's collaboration components allowed experts to monitor user consultations remotely, view summaries of responses, and trace-rule inference chains. Experts and users participated in real-time chat sessions or posted questions on extended discussion lists. The system allowed experts and users to experiment with real-time enhancements of knowledge bases. Fuzzy rules resulted in semantically richer knowledge bases that flexibly handled complex and uncertain knowledge. A fuzzy inference engine supported hedges and partial matching to assist users in applying knowledge and exploring Web-based data.

Keywords: expert system, fuzzy logic, Internet, collaboration, design

ACM Categories: H.4.2, I.2.1, I.2.5

Introduction: Challenging Expert System Assumptions

Expert systems attempt to clone human expertise to avoid geographical and time-based limitations of consulting with human experts. Turban and Aronson (1998, p. 17) present the main idea behind expert systems: "Expertise is transferred from the expert to the computer...users call on the computer for specific advice as needed." The expert's role is to assist knowledge engineers in developing, refining, and testing the knowledge base. Once the knowledge base is delivered, the expert is not involved in assisting specific users, and users are not involved in maintaining the knowledge base.

This paper explores a convergence of Internet and fuzzy logic technologies that are challenging these assumptions. Internet technology removes geographical and time-based constraints to improve users' access to human expertise. With Web-based expert systems, experts and users can more easily collaborate in problem solving and knowledge base development. Fuzzy logic technologies reduce cognitive dissonance by coding knowledge in English-like expressions which provide an opportunity to better involve users in maintaining knowledge bases. Compared to standard-rule representations, the numbers and complexity of rules required in fuzzy models are much less than with traditional knowledge bases (Cox, 1995).

The following scenario demonstrates the potential for improved communication and ongoing maintenance of globally distributed knowledge bases.

As the user, you are responsible for troubleshooting printer problems on your network. You tie into a help-desk site with your Web browser that feeds you a Java fuzzy expert system. You select a knowledge base for your printer problems and start the consultation. During the consultation the system requires that you evaluate a document's print quality. You specify that the print quality is between poor and average. Unfortunately, the consultation's recommendations prove inadequate to solve your problem and you request an on-line chat with a human expert. You are in luck because the system is undergoing a maintenance review, and the experts are currently available to monitor consultations. You take your place in a queue, and an audible alarm alerts you that an expert is available. The expert remotely observes as you repeat your previous consultation. During the consultation you and the expert review the fuzzy rules concerning print quality. A rule's fuzzy premise states — *if print quality is **significantly** poor.* You have an on-line chat with the expert describing your print quality. The expert agrees to change temporarily the premise of the rule to — if *print quality is **somewhat** poor.* Another consultation session with the revised rule then resolves the printer problem. The expert's change, and the consultation results are logged for the knowledge base's next maintenance review.

This paper explores architecture for supplementing knowledge-based expert systems with Internet technology and fuzzy models. The following sections review research issues, describe a system architecture for development and delivery, discuss our experiences with prototype consultations delivered over the Web, and conclude and raise future research possibilities.

Research Issues

Over 12,500 expert systems are deployed in manufacturing, medicine, and business; and the number of experts system development tools has been growing at about 16% per year (Durkin, 1996). Nevertheless, expert system developers continue to struggle with design issues such as knowledge acquisition, testing, and maintenance. The Internet and local intranets offer new ways to deliver knowledge-based advice. Traditional shells, however, do not support openness and interoperability required for deploying expert systems over wide-area networks. In addition, globally accessible knowledge bases are difficult to maintain and update. Researchers are exploring new Web-based applications of expert systems, and the following highlights their current efforts.

One technique for deploying experts systems on the Internet relies on Common Gateway Interfaces (CGI) to coordinate client/server interactions (MultiLogic Inc., 1998; Inference Corp., 1996; Bello and Ribeiro, 1995). The CGI server is responsible for controlling logical inferences, maintaining the system's knowledge, and dynamically constructing and distributing HTML forms to conduct consultations. Java technology offers a new way to deliver expert systems where the user's browser serves as an interface for the consultation applet (Ernest Friedman-Hill, 1996). Java-based shells deliver knowledge and provide automatic access to Internet resources but otherwise structure consultations in the same manner as traditional expert systems.

Autonomous software agents provide another form for Internet delivery of knowledge bases. Software agents cooperate to exchange knowledge to solve user problems. Researchers are focusing on agent communication structures and interoperability to improve cooperation and support networks of distributed knowledge bases (Genesereth and Ketchpel, 1994).

In contrast to improving communications among software agents, others are focusing on applying the Internet to improve interpersonal communications for knowledge acquisition. Shaw and Gaines (1997) have developed REPGRID — an interactive Web-based system to elicit personal constructs for collaborative learning. The system supports knowledge acquisition through interactive repertory grid analysis, where knowledge engineers and experts collaborate from different geographical locations. Others have proposed that knowledge-based systems on the Internet can support collaboration efforts by assisting in constructing user interfaces and supporting system designs (Benford et al. 1993; Nakakoji, and Fischer, 1995; Far and Koono, 1996).

The Internet offers an opportunity to combine the synergies of expert systems and groupware technology. Expert system tools improve groupware by assisting in information retrieval, simplifying system operations, and structuring interactions (Aiken et al. 1991). Groupware tools support expert systems by promoting collaboration and maintaining distributed knowledge bases.

Maintaining and enhancing Web-based expert systems require a method to help users and designers manage uncertainty caused by lack of available data, ambiguous classification categories, and semantic misunderstandings (Cox, 1995; Walz, Elam and Curtis, 1993). End users, designers, and experts must make judgments, assess the appropriateness for applying rules, and share responsibly for the safety and legal implications that result from misdiagnosis or misclassifications (Mira, Yanez, and Barreiro, 1991). Fuzzy models respond to these concerns by providing a means to represent and adapt to inherent vagueness and ambiguities that can result when applying a general rule to a specific situation (Klir and Folger, 1988).

Fuzzy expert systems reason through fuzzy logic membership functions. Membership refers to the degree to which a particular attribute's value belongs to a set. For example, someone with an age attribute equal to 21 years would have high membership in the set 'young' and a low membership in the set "old." Membership functions allow degrees of membership to be related to linguistic terms (Zadeh, 1965).

Fuzzy logic in expert systems allows fuzzy propositions of the form: if size is more or less *small* then investment is rather *large,* where *small* and *large* are linguistic variables denoting fuzzy memberships and more or less and rather are hedges that modify memberships. Fuzzy expert systems also apply fuzzy numbers representing degrees of membership over intervals. Table 1 defines fuzzy variables, hedge and fuzzy numbers.

Fuzzy expert systems allow partial matching of a rule's antecedents to provide a systematic way of managing imprecision and uncertainty. Compared to traditional expert systems, fuzzy expert systems take less time to develop, reduce maintenance cost, and improve user understanding (Cox, 1995; Schneider et al. 1996).

As the amounts and diversity of information in databases and on the Internet continue their exponential growth, users are challenged to locate data resources and create queries that are aligned to support specific decisions. Internet retrieval tools such as search engines, query-by-form, and intelligent "wizards" provide assistance in the mechanics of information retrieval but offer limited support for helping users comprehend a database's context and contents (George, Buckles, Petry, & Radhakrishnan, 1996). A knowledge base combined with a fuzzy query system helps users obtain knowledge-based assistance to form queries in commonly understood terms.

Internet-aware knowledge bases improve communication and productivity for a variety of information systems. Figure 1 shows the diversity of IS environments requiring support where knowledge bases can assist in database connectivity and offer mobile agents — "traveling experts" — to assist in accessing enterprise and middleware resources across wide-area networks.

This review suggests that combining fuzzy expert systems and groupware tools improve knowledge development, delivery, and maintenance. Fuzzy techniques result in knowledge bases that are flexible and semantically rich. Internet-based groupware allows experts and end users to collaborate in managing and accessing knowledge and data resources. The following describes a prototype development and delivery environment to investigate collaborative fuzzy expert systems delivered on the Web.

Instant Fuzzy Traveling Expert Advice

ITEA (Instant Traveling Expert Advice) is both an integrated development environment (IDE) application and Java delivery applet for producing, delivering, and collaborating to maintain fuzzy knowledge bases on the Web. The entire system is implemented in Java to promote portability and ease of networking. Figure 2 shows the system's components.

The delivery applet conducts user consultations by applying knowledge bases created by the IDE. Consultations take place within a Web browser, where users respond to system generated questions. Based on the user's responses and fuzzy logical inferences, the system retrieves data, derives problem solutions, and presents recommendations. The following describes system components and presents examples of fuzzy rules.

Fuzzy Variable	A fuzzy variable is an attribute described by a set of linguistic values. Sets of values are related by a membership function to express possibilities of set membership for particular instances. For example, terms representing increasing membership such as narrow, medium, and wide may describe a shoe's broadness attribute. An instance of a shoe may be described as medium broadness, indicating that shoe's membership in "broadness" is intermediate between narrow and wide.
Hedges	A hedge is a linguistic term to qualify fuzzy variables by increasing, reducing or restricting membership levels. For example, hedges such as "very," "strongly," and "really" denote increased membership. Hedges such as "more or less," "slightly," and "few" denote decreased membership. Hedges such as "relatively," "technically," and "strictly" denote restrictions on membership levels. A shoe with **slight** medium broadness indicates reduced membership compared to a shoe with medium broadness (Bouchon-Meunier,1992).
Fuzzy Numbers	A fuzzy number relates numeric intervals to degrees of possibility associated with a proposition. For example, a fuzzy number may relate a temperature level to the proposition that the temperature is warm. A temperature of 32° F. or (0° C.) would have low membership in warm. While 60° F. would have a higher membership. Two common forms of fuzzy numbers are triangular and trapezoidal.

Table 1. Definitions of fuzzy terms.

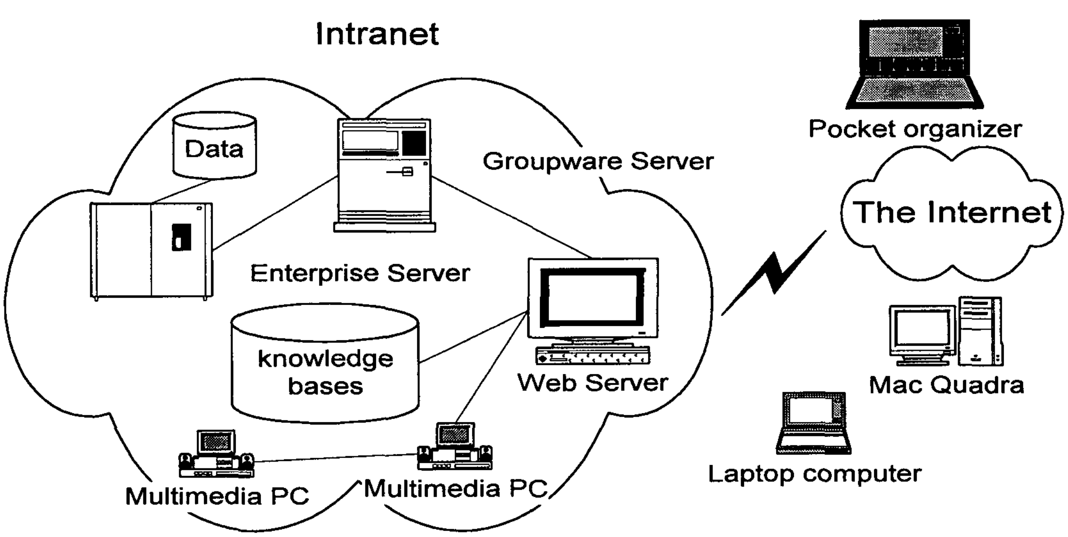

Figure 1. Knowledge bases can support database retrieval, and decision making across Intranets and the Internet.

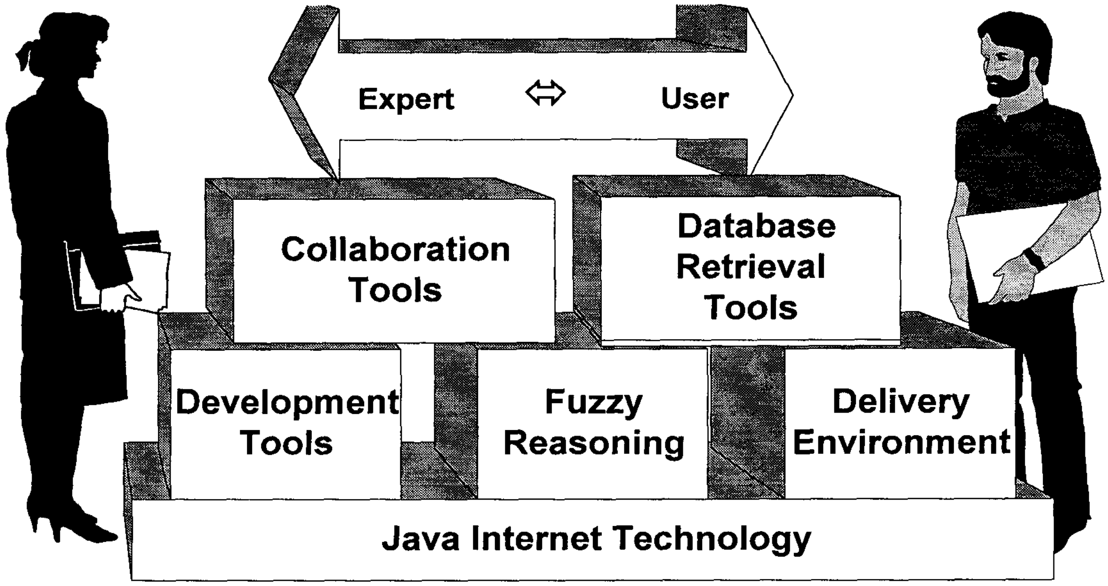

Figure 2. System components for ITEA.

The Development Environment

The Java IDE features provide an object-oriented editing environment for creating classes, subclasses, attributes, and values. Developers create IF-THEN rules that support logical operations ("AND," "OR," "NOT," "ELSE"), fuzzy numbers, fuzzy variables, and fuzzy reasoning. The system provides a fuzzy membership function generator that supports hedges for fuzzy linguistic variables and allows developer to specify triangular or trapezoidal memberships functions for fuzzy numbers. Developers set the system's inference strategy to pursue single or multiple values for goals and specify confidence thresholds for firing rules.

The Delivery Environment

Figure 3 displays the applet interface for the delivery environment. The delivery environment consists of libraries of Java classes that can be deliveredacross any platform. The expert system is a Java applet that is executed within a user's browser. Users may select from a list of knowledge bases residing at a Web site. The system conducts the consul-tation and may retrieve additional Web-based resour-ces to present questions, provide explanations, explore databases, and retrieve solutions. At the conclusion of a consultation, the user may review the session's rule inference chains, repeat the consultation, or select another knowledge base.

Collaboration and Explanation Environment

The delivery of expert system consultations on the Web provides an opportunity to improve interaction between experts and users. Figure 4 displays the system's interface for remote collaboration. The collaboration components allow experts to monitor users consultations remotely, view summaries of users responses, and trace inference chains. Further, experts and users can participate in either real-time chat sessions or post questions on extended discussion lists. During real-time chat, experts make on-the-fly changes to rule bases to allow users to remotely experiment with knowledge base changes. Knowledge base changes and enhancements are maintained at a central site and automatically distributed user sites.

Reasoning with Fuzzy Variables

The following demonstrates how ITEA manages inferences with fuzzy variables. For instance, consider the following rules:

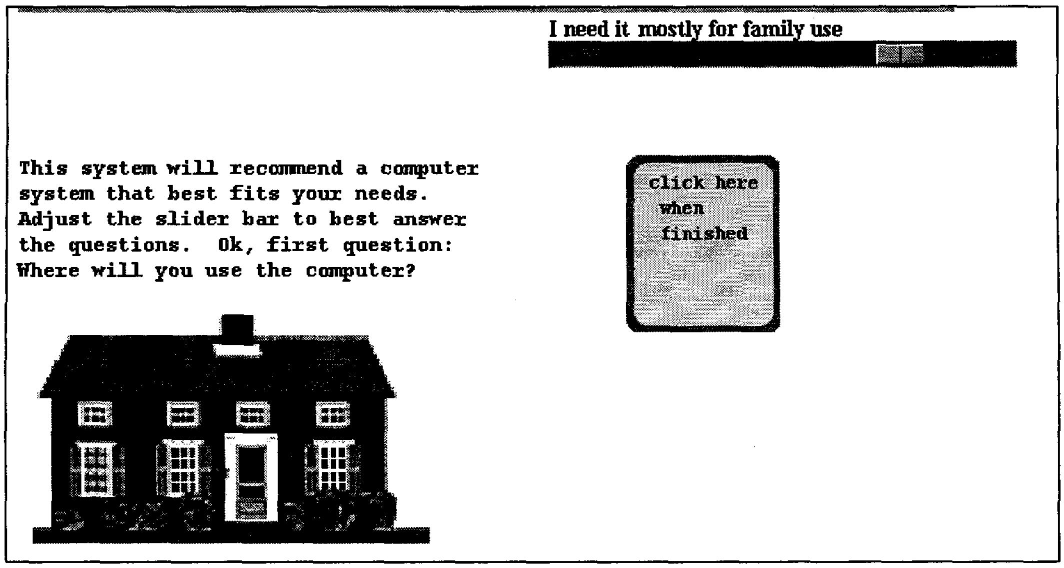

Figure 3. User interface for the delivery environment.

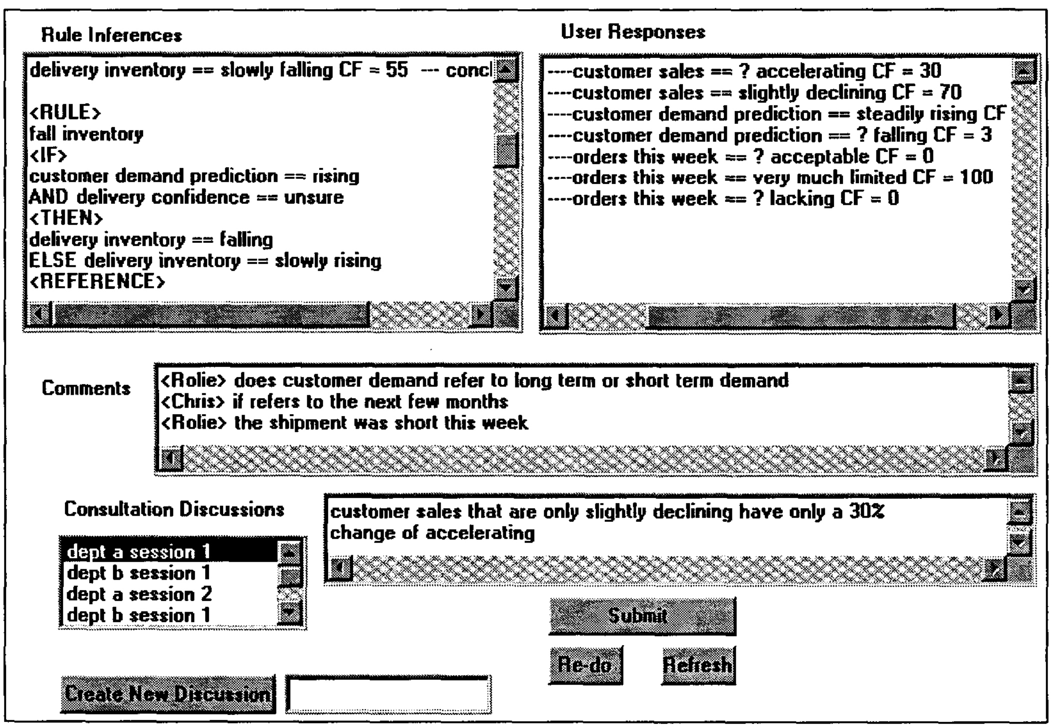

Figure 4. Collaboration tools for remote session monitoring and chat.

Rule 1:
If environment humidity == more or less low
AND environment wind == breezy
AND environment temperature == mostly warm
THEN weather condition == more or less nice

Rule 2:
IF weather condition == slightly nice
AND snow condition == mostly good
OR snow condition == very ok
THEN recommend activity == ski

Assume the following fuzzy attributes have been defined for application classes - **environment**, **weather** and **snow**. The attributes for **environment** are:
humidity with fuzzy values — (high, medium, low) and hedges (very, more of less, slightly)
wind with fuzzy values — (breezy, calm) and hedges (strongly, slightly)
temperature with fuzzy values — (hot, warm, cold) and hedges (mostly, not very)

The attributes for **weather** are:
condition with fuzzy values — (the best, great, nice, ok, terrible, the worst) and hedges (very, more or less, slightly)

The attributes for **snow** are:
condition with fuzzy values — (excellent, good, ok, bad, awful) and hedges (very, mostly, slightly)

The goal of the system is to provide a recommended activity. Consider the fuzzy inferences for a consultation where the user provides the following fuzzy facts:
humidity is very low
wind slightly breezy
temperature is mostly warm
snow condition is very good

The system reasons by backward chaining to assign fuzzy memberships that depend on the user's responses and the possibility distributions defined by a rule's hedges and fuzzy attributes. For the responses above, the following can be concluded:

The first premise statement of Rule 1 is true to a fuzzy certainty level of 85%, the minimum membership attained among the first rule's three premise statements:

IF environment humidity == more or less low
AND environment wind == breezy
AND environment temperature == mostly warm

The fuzzy certainty level was determined by the following procedure:

Given the user response that humidity is very low, a membership of 100% is assigned to the Rule 1 first proposition requiring humidity to be more or less low.

The user response that the wind is slightly breezy allows the system to assert that wind has an 85% membership in breezy and a 15% membership in calm. The rule's statement contains no hedges so that membership assignment is not changed.

The user response that the temperature is mostly warm warrants that the third statement is assigned a 100% membership in mostly warm.

Since the Rule 1 premise's fuzzy certainty of 85% exceeds the system's threshold level of 50%, Rule 1 fires to conclude the weather condition is nice. The rule fires by adding the fact that the weather is nice to the system's working knowledge and adjusting the confidence according to the conclusion's hedge. Nice at a membership level of 85% (assigned by the premise) warrants a 100% membership for the hedged fuzzy set more or less nice (see Figure 5). So, the system asserts that weather is more or less nice adjusted to a membership of 100%.

Rule 2 assigns the attribute — recommended activity with the values — ski with a confidence level of 100% for the following reasons:

Rule 2 premises
IF weather condition == slightly nice
AND snow condition == mostly good
OR snow condition == very ok

More or less nice weather is known by the system to be true at a level of 100% from Rule 1. Since the first statement of the Rule 2 premise requires only a weather condition of slightly nice, this statement results in an assignment of 100% membership for slightly nice. That is, a weather condition that is more or less nice is also slightly nice.

The user's assessment that the snow condition is very good results in assigning a 100% fuzzy certainty for the premise's second statement, where the statement requires a slightly good snow condition.

The rule membership assignment depends on both the user responses and the construction of a rule's statements.

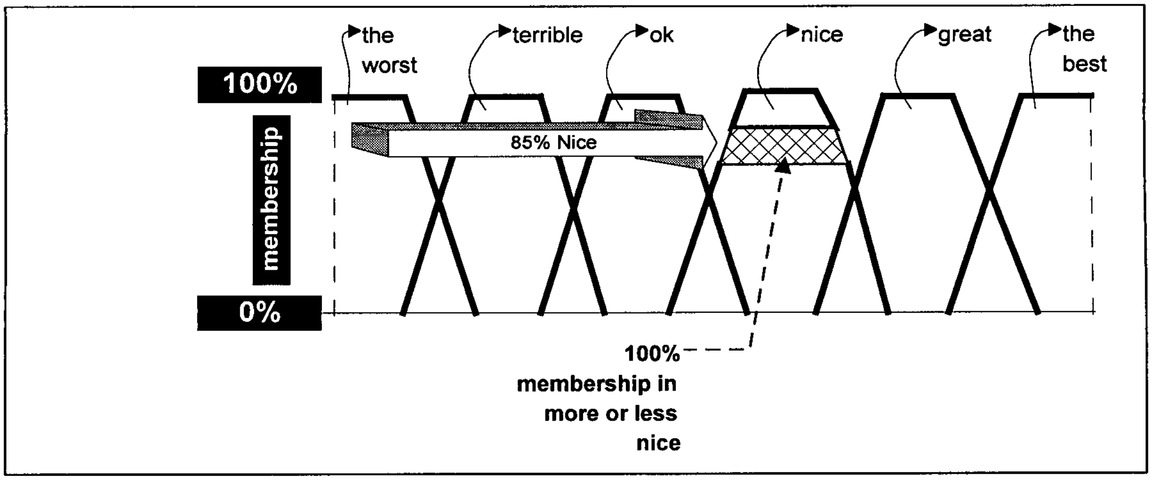

Figure 5. Membership functions for fuzzy values of attribute weather. A fuzzy membership of 85% for the fuzzy value *nice* translates in to 100% membership for more or less *nice*.

If the user had answered that the snow condition was mostly ok, the rule would fail since Rule 2 requires the snow to be at least very ok.

If the first statement in the Rule 2 premise had required the weather condition to be very nice, the rule would fail since the Rule 1 conclusion warrants only that the weather condition is more or less nice.

Database Retrieval Module

ITEA supports exploration of Internet databases by providing assistance in defining information needs, determining relevant information sources, and formulating fuzzy queries. The knowledge base controls database retrieval by dynamically constructing initial queries, retrieving data, calculating fuzzy membership, and presenting query results.

The system's knowledge base applies an expert's domain knowledge to define a retrieval vocabulary that reflects qualitative distinctions within a domain. For example, in a medical domain, physicians distinguish the onset of a disease by fuzzy terms such as acute, rapid, or normal. Users are better prepared for future collaboration as they explore and learn about complex relations by applying a domain's terminology (Larsen and Yager, 1997; Xu and Ichikawa, 1992).

The knowledge base pre-classifies users and recommends query profiles that are likely to satisfy a specific category of users. For example, users may be classified according to level of experience, where beginners receive detailed guidance in formulating queries. Experienced users may dispense with consultation and instead directly explore databases.

The knowledge base communicates with a Fuzzy Database Retrieval Module to recommend data sources, vocabulary, and attribute sets for constructing queries, and explanations. The Retrieval Module is then responsible for connecting with Internet databases and retrieving information. The user may then refine the query by selecting different combinations of attributes, applying different query operators, or raising and lowering acceptable thresholds. After database exploration, the user returns to the consultation session where the knowledge base's explanation model summarizes items retrieved and may provide suggestions for future explorations.

Discussion

This section describes our experiences and discusses the application of ITEA for delivery and maintenance of Web knowledge bases. The following describes a prototype application that helps users to rank sets of banks according to financial ratios.

Financial analysts value banking stocks through common financial ratios, such as return on investment, price-earnings, price-book, and other ratios unique to the banking industry such as net interest margin and efficiency ratio. Net interest margin measures the spread between interest rates for loans and deposits, divided by total earning assets (loans plus securities). The efficiency ratio measures the proportion of operating expense to revenue.

The knowledge base helps users form queries that express financial ratios through fuzzy terms instead of quantitative limits. For example, financial analysts consider a bank well managed if its efficiency ratio is low, and they suggest that 60% is a good target. The analysts also know that a price-earnings (P/E) ratio of 21 is good and that 24 is very good (Schlegel, 1997). The analyst may then perform queries and implicitly rank stocks based on their expertise in construing financial ratios.

A fuzzy query, in contrast, requires only that novice users specify banks with "low" efficiency ratios and "very good" P/E ratios. The system is responsible for assigning fuzzy memberships to banks through fuzzy comparisons within a group of bank stocks. The knowledge base assists users to formulate queries initially by explaining financial ratios, suggesting preformulated queries that best match a user's investment style, and guiding query construction and automating retrieval.

The knowledge base classifies users according to risk and investment goals. Conservative investors are willing to accept lower returns in exchange for safer investments. Moderate investors seek moderation in returns and risks, while aggressive investors are willing to risk losses in exchange for the possibility of larger returns. The knowledge base contains rule-based advice for selecting banks based on a user's profile. The following rule defines a preformulated fuzzy query for aggressive investors:

```
<RULE>
aggressive investor
<IF>
investor desire == match a query to my style
AND investor preferences == aggressive
<THEN>
query! price == none
AND query! price-book == high ratio
AND query! price-earnings == high ratio
AND query! return on equity == very high ROE
AND query! net interest margin == very wide
AND query! efficiency ratio == very low
```

The knowledge base also offers guidance to help users interactively construct fuzzy queries. For example, during query formulation the following advice is offered:

"A price-to-earnings ratio is the share price divided by earnings per share for the company's most recent four quarters. A very high P/E ratio may indicate an overpriced stock. But remember there may be a good reason for the premium in stock price. You probably should select a lower ratio for conservative stocks."

Users may decide to follow the knowledge base's recommendation or formulate customized queries. The system applies the user's responses to define a query and locate Internet databases and communicate results.

The knowledge base dynamically constructs an initial query, connects with database resources, offers query refinement, and presents results. Figure 6 presents the interface where a user reviews query results and formulates fuzzy queries to further explore bank stocks.

At the completion of user explorations, query results are returned to the knowledge base for further analysis. The knowledge base provides additional information concerning individual banks and suggests other Internet resources to further evaluate companies.

The initial evaluations of the prototype focused on user interaction and understanding. An exploratory evaluation suggests that users' final investment decisions and explanations were consistent with expert recommendations for selecting bank stocks. Results indicate that users valued support for constructing queries and easily explored databases with fuzzy queries. Specifically, the tools capture and apply expertise to assist novices in aligning queries with user preferences, promoting understanding of the information contained in a domain, and allowing users to discover data patterns interactively and better interpret results.

Collaboration Strategy

The collaboration subsystem was designed to improve user interactions and assist in knowledge-base development and maintenance. The computer selection knowledge base was placed on a Web server January 1997, and since then there have been thousands of consultations sessions from users representing numerous countries. The following summa-

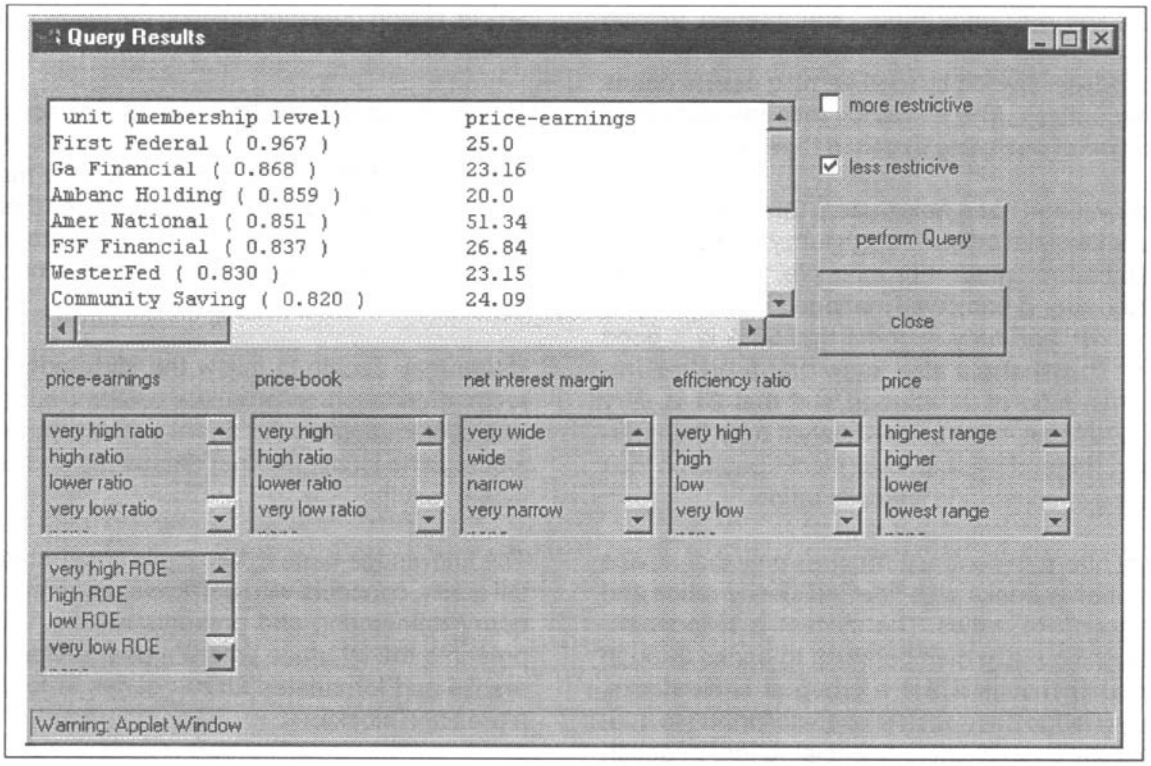

Figure 6. Fuzzy query results.

rizes early experiences with the collaboration subsystem.

Users found it easy to provide feedback, and their suggestions included adding additional rules for Macintosh and other computer brands, adding price considerations, reworking question syntax, providing new Internet resources links, and modifying hedges in rules (e.g., changing moderately unimportant to slightly important). Other users alerted designers to interface and display problems associated with running the system on their platform and suggested new features such as improved graphic handling and enhanced reasoning facilities.

We have recently added on-line session monitoring and real-time chat to increase levels of collaboration between users and experts. Initial experiences suggest that users are eager to share real-world experiences to assist knowledge designers and experts. Collaborations allowed experts to experiment by modifying and adding rules to respond to user concerns immediately. Experimental versions of the knowledge bases were then maintained at a central site along with the established version. Other user sites could then access the experimental versions and provide comments.

The combination of expert systems delivery and groupware addressed real-time operational and maintenance concerns. Experts provided answers to user questions to assist users during consultations. Users provided designers with insight into user problem-solving approaches. The systems collaboration features offered an opportunity to increase user participation, enhance the knowledge base, and centralize distributions of updates.

Currently there are several ITEA knowledge bases deployed at several locations on the Web including (see http://www.instanttea.com for demonstration systems):
Movie advisor — recommends a movie and provides access to resources
Sports advisor — an intelligent index to sports related internet sites
Fish disease diagnosis — a diagnosis system with links to remedies
Mutual fund advisor — a classification system for mutual funds
Snake identification — identifies regional snakes
Conference planning — offers intelligent indexes to conference information
Car diagnosis — a troubleshooting system
Golf club selection — recommends clubs for vari-

ous course situations
Computer selection — recommends computer brands with links to suppliers

These systems are available on the Web 24 hours a day and take advantage of distributed Internet resources to present explanations and conclusions. ITEA's delivery and development environment is freely available to researchers and educators by contacting the author.

Conclusion

The convergence of technologies including Java, the Internet, and fuzzy logic provides a new solution for delivery and maintenance of expert systems. Internet technology allows knowledge to be delivered anywhere in the world. Fuzzy logic supports construction of rule bases that are easier to understand and maintain.

This paper presented ITEA as a Web-based fuzzy expert system to promote collaboration and improve expert system delivery and maintenance. Web-based expert systems allow users easy access to Internet resources and provide designers a way to maintain and distribute knowledge from a central location. Fuzzy rules result in semantically richer knowledge bases that flexibly handle complex and uncertain knowledge. Integrated collaboration features that support real-time chat and rule modifications show promise for improving knowledge base usability and rule maintenance.

The first prototype of ITEA was developed in the fall of 1996. The collaboration components were added in early 1997. We are currently prototyping three-tier modules, object-oriented databases, and new protocols to improve delivery and collaboration.

Our plans include testing usability among experts and users to evaluate how ITEA meets knowledge engineering goals, refining an expert system development methodology that provides increased levels of ongoing collaboration between end users and experts, and integrating and investigating Web knowledge bases and databases.

The integration of the Web resources with the field of fuzzy expert systems offers new ways of sharing and distributing knowledge in an organization. Users and experts can jointly access corporate resources to explore fuzzy models of complex situations collaboratively. This also offers an opportunity for simulating and refining the modeling of business processes (Yu et al, 1996; Fishwick, 1991). A collaborative fuzzy expert system for the Web supports a new design for distributed decision making.

References

Aiken, M.W., Sheng, O, and Vogel, D. (1991). "Integrating Expert Systems with Group Support Systems," *ACM Transactions on Information Systems*, Vol. 9, No. 1, January, pp. 75-95.

Benford, Steve, Bullock, Adrian, and Harvey, P., and Grace, P. (1993). "A System to Support the Development and Use of Global Computer-Supported Cooperative Work (CSCW) Applications," *Internet Research*, Vol. 3, No. 1, Spring, p. 25.

Belo, O., and Ribeiro, A. (1995). "A Web-Based Framework for Distributed Expert Systems," *Conferência Nacional WWW* (CNW3'95), Universidade do Minho, Braga, Portugal, Julho.

Bouchon-Meunier, B. (1992). "Fuzzy Logic and Knowledge Representation Using Linguistic Modifiers." *Fuzzy Logic for the Management of Uncertainty*," (Zadeh and Kacprzyk, eds.), New York, NY: Wiley.

Cox, E. (1995). *Fuzzy Logic for Business and Industry.* Charles River Media. Rockland, MA.

Durkin, J. (1996). "Expert Systems - A View of the Field," *IEEE Expert/Intelligent Systems and Their Applications*, Vol. 11, No. 2, April, pp. 56-63.

Far, B.H., and Koono, Z. (1996). "Ex-W-Pert System: A Web-Based Distributed Expert System for Groupware Design," *World Congress on Expert Systems ' 96*, Seoul, Korea, February, pp. 545-552.

Fishwick, P. (1991). "Fuzzy Simulation: Specifying and Identifying Qualitative Models," *Int. J. General Systems*, Vol. 19, pp. 295-316.

Friedman-Hill, Ernest and Sandia National Laboratories (1996). *Jess Java Expert System Shell.* URL:http://herzberg.ca.sandia.gov/jess. MultiLogic Inc. (1998). URL: http://www.multilogic.com.

Genesereth, M., and Ketchpel, S. (1994). "Software Agents," *Communications of the ACM*, Vol. 37, No. 7, July, pp. 72-76.

Inference Corporation's Case-Base Reasoning Search Engine on the Web (1996). URL: http://m5.inference.com/wcp/.

Klir, G.J., and Folger, T.A. (1988). *Fuzzy Sets, Uncertainty, and Information*, Englewood Cliffs, NJ: Prentice-Hall.

Larsen, H.L., and Yager, R. (1997). "Query Fuzzification for Internet Information Retrieval." *Fuzzy Information Engineering* (eds., Dubois,

Prade, & Yager). New York: John Wiley & Sons, Inc., pp. 253-264.
Mira, J., Yanez, A., and Barreiro, A. (1991). "Fuzzy Specification in Data Capture and Knowledge Representation For an Expert System in Oncology," *Fuzzy Sets and Systems*. Vol. 44, No. 3, December, pp. 431-448.
Nakakoji, K., and Fischer, G. (1995). "Intertwining Knowledge Delivery and Elicitation: a Process Model for Human-computer Collaboration in Design," *Knowledge-Based Systems.*, Vol. 8, Nos. 2-3, April, pp. 94-106.
Shaw, M., and Gaines, B.R. (1997). "Comparing Constructions Through the Web," *Knowledge Science Institute,* University of Calgary, Calgary, Alberta, Canada T2N 1N4h, ftp://ksi.cpsc.ucalgary.ca/articles/CSCL95WG/.
Schneider, M., Kandel, A., Langholz, G., and Chew, G. (1996). *Fuzzy Expert System Tools*, New York, NY: Wiley.
Turban, E., and Aronson, J. (1998). *Decision Support Systems and Intelligent Systems*, Englewood Cliffs, NJ: Prentice Hall.
Walz D.B., Elam, J., and Curtis, B. (1993). "Inside a Software Design Team: Knowledge Acquisition, Sharing and Integration," *Communication of the ACM.* Vol. 36, No. 10, October, pp. 63-76.
Xu, Wu, and Ichikawa, Tadao. (1992). "KDA: A Knowledge-based Database Assistant with a Query Guiding Facility." *IEEE Transactions on Knowledge and Data Engineering*, Vol 4, No. 5, pp. 443-453.
Yu, E., Mylopolulos, J., and Lesperance, Y. (1996). "AI Models for Business Process Reengineering," *IEEE Expert/Intelligent Systems and their Applications.* Vol. 11, No. 4, August, pp. 16-23.
Zadeh, L.A. (1965). "Fuzzy Sets," *Information and Control*, Vol. 8, pp. 338-353.

About the Author

Tod A. Sedbrook is an associate professor at the College of Business Administration, University of Northern Colorado, Greeley, Colorado. He received his Ph.D. in management information systems from the University of Colorado. His research interests include Internet delivery of instructional resources and methodologies of object-oriented analysis and design. He has published in the areas of knowledge-based systems that apply fuzzy logic, genetic algorithms, neural networks and statistical techniques. He is the editor of the *International Business Schools Computing Quarterly* and President of Instant Traveling Expert Advice, Inc.
E-mail: tasedbr@unco.edu