ed papers on case-based guidance for hypermedia with cases indexed by neural net and with a payoff matrix model to determine optimal actions (including communications) for a decision-support system. The session on natural language featured talks on selecting dialog moves by using a subordinate instantiation of the system (global anticipation feedback) to predict user responses, dialog strategies for tutoring the user (modeled as a weighted overlay) about causality in physics; dividing up explanations to avoid cognitive overload based on a parametric process model of concept learning; and generating arguments in the chess domain by selecting paths through an argument graph based on a model of the user's beliefs, inferences (sound and unsound), and ability level. The session on cognitive models consisted of reports on a nonmonotonic belief logic for formalizing reasoning about the beliefs of incomplete agents (that is, agents that sometimes overlook inferences), experiments that show users' self-confidence (their own performance prediction) does not match actual performance but tracks the difficulty of the hardest elements of the task, and a user model for modeling fault diagnosis that combines an overlay of the system knowledge with a mean-ends analysis of the diagnosis process.

The UM-96 program cochairs were Sandra Carberry (University of Delaware) and Ingrid Zukerman (Monash University, Australia). The local arrangements chair was Martha Crosby (University of Hawaii), and the conference chair was David N. Chin (University of Hawaii). We are indebted to the American Association for Artificial Intelligence for sponsorship of student travel grants, Kluwer Academic Publishers for sponsoring the best paper award, and the University of Hawaii for equipment loans.
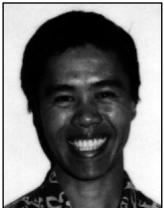
David N. Chin is an associate professor of information and computer sciences at the University of Hawaii. He has been heavily involved in establishing and promoting the field of user modeling. His e-mail address is chin@hawaii.edu.

AAAI-96 Workshop on Internet-Based Information Systems

Alexander Franz

The Workshop on Internet-Based Information Systems, held as part of the Thirteenth National Conference on Artificial Intelligence, provided a venue for researchers to discuss the role of AI and related fields in network-based information systems. Not surprisingly, most of the discussion focused on the World Wide Web (WWW), and a good part of the workshop was devoted to topics related to WWW search and retrieval. A number of presentations addressed ways to improve search accuracy and efficiency. Methods under discussion included statistical modeling, machine-learning techniques, and the addition of ontologically categorized metadata to WWW pages. An invited talk by Michael Mauldin of Lycos Inc., who described the development, commercial deployment, and current state of the LYCOS search engine, provided an industrial perspective on this topic. It seems certain that research in this area will continue to expand. At the same time, the successful transfer of research results to commercial search engines will require persuasive cost-benefit analyses for new proposed methods.

A second major theme was the use of natural language–processing techniques for analyzing WWW pages. For example, two presentations described means of taking advantage of the question-answer format of frequently asked question (FAQ) files to provide more precise information-retrieval abilities (albeit limited to the narrowly circumscribed domains covered by FAQ files). To address the bottleneck of manually creating hypertext from linear documents, another presentation described a technique for the automatic creation of hypertext links between related passages.

A second invited talk by Peter Pirolli of Xerox Parc described an ecologically inspired model, based on optimal foraging theory, that views WWW browsing as foraging for information. A related theme that emerged at the workshop concerns the study of the differences between various information services on the internet, such as talk, internet relay chat, e-mail, mailing lists, news groups, and web pages. We can expect further research into human information-seeking behavior on the internet, especially because accurate predictive models could result in immediate practical benefits, such as improved caching and prefetching schemes for HTTP (hypertext transfer protocol) proxy servers. On the whole, the considerable interest in this workshop, as well as other activities on internet-related topics at AAAI-96 and elsewhere, makes it evident that the AI research community's interest in the internet continues to grow apace. Further information about the workshop is located at www.csl.sony.co.jp/person/amf/iis96.html.

Alexander Franz is a researcher at Sony's Computer Science and D21 Laboratories. He received a Ph.D. in computational linguistics from Carnegie Mellon University. His research interests include speech, natural language, and multilingual information and communication. His e-mail address is amf@csl.sony.co.jp.