

AAAI 1997 Spring Symposium Series Reports

■ The American Association for Artificial Intelligence (AAAI) held its 1997 Spring Symposium Series on 24 to 26 March at Stanford University in Stanford, California. This article contains summaries of the seven symposia that were conducted: (1) Artificial Intelligence in Knowledge Management; (2) Computational Models for Mixed-Initiative Interaction; (3) Cross-Language Text and Speech Retrieval; (4) Intelligent Integration and Use of Text, Image, Video, and Audio Corpora; (5) Natural Language Processing for the World Wide Web; (6) Ontological Engineering; and (7) Qualitative Preferences in Deliberation and Practical Reasoning.

Artificial Intelligence in Knowledge Management

Knowledge management (KM) is a topic of growing interest to large organizations. It comprises activities focused on the organization, acquiring knowledge from many sources, including its own experience and that of others, and on the effective application of this knowledge to fulfill the mission of the organization.

The KM community has been eclectic in drawing from many sources for its methodologies and tools. Typical approaches to the management of knowledge are based on concept maps, hypermedia, and object-oriented databases. Techniques developed in AI for knowledge acquisition, representation, and discovery are seen as relevant to KM. However, as yet, KM has no unified underlying theory, and the scale of the problem in large organizations is such that most existing AI tools cannot be applied in their current implementations.

The objective of this symposium was to bring together KM practitioners and applied AI specialists from knowledge acquisition, knowledge representation, and knowledge dis-

covery in databases and attempt to formulate the potential role of various AI subdisciplines in KM. Those attending represented a wide range of industries and areas relevant to AI research and application.

The symposium began with keynote addresses on industrial requirements for KM by Vince Barabba of General Motors and Rob van der Spek of CIBIT. The remainder of the meeting was devoted to intensive group discussion of the role of AI in KM, including a joint session with the Symposium on Ontological Engineering. The sessions were entitled Organizational Knowledge Management, Work-Flow Management, Knowledge Management through Hypermedia and Text Processing, Cognitive Aspects of Knowledge Management, Agents and Multiactor Systems, Knowledge Representation and Reasoning, and Knowledge Discovery.

The accepted papers were made available to participants in advance of the meeting through the World Wide Web and will remain available at ksi.cpsc.ucalgary.ca/AIKM97.

B. R. Gaines

University of Calgary, Canada

Mark A. Musen

Stanford University

Ramasamy Uthurusamy

General Motors

Computational Models for Mixed-Initiative Interaction

This symposium was held as a direct result of discussions at the National Science Foundation Interactive Systems Program Grantees Workshop that was held in Cambridge, Massachusetts, in November 1995. Mixed-initiative interaction is a

notion that has concerned members of the interactive systems community for some time, although there is no established vocabulary for discussing it and no established method for achieving it. An important goal of this symposium was to define *initiative* and consider examples of how systems currently support it.

The symposium brought together researchers in natural language processing, planning, robotics, interactive tutoring systems, and human-computer interfaces. During the symposium, discussion transitioned from theoretical issues to demonstrations of working systems that exhibit varying degrees of initiative. On the first day, participants attempted to define concepts such as *having control* or *taking the initiative*. Participants also addressed practical questions such as whether domain task models include enough information to reason about initiative and control or whether discourse-level information is needed. On the second day, participants examined the role that mechanisms such as scripts, constraint satisfaction, blackboards, and top-down planning can play in monitoring and managing an interaction. On the third day, participants watched and discussed demonstrations of interactive systems.

Participants concluded that (1) mixed-initiative behavior is required in any system that aims to tutor users or help them solve problems collaboratively and (2) the problem of getting systems to demonstrate mixed-initiative interaction is a serious one. For some, achieving initiative requires including explicit reasoning about initiative as part of the design, but for others, it can be (and is) an emergent property of the architecture. To better understand the problem, participants will be collecting and annotating interactions to identify changes in control and initiative.

Susan Haller

University of Wisconsin at
Parkside

Susan McRoy

University of Wisconsin at
Milwaukee

Cross-Language Text and Speech Retrieval

Cross-language text retrieval (CLTR) is the problem of matching a query in one language to related documents in other languages. As internet resources such as the World Wide Web have become global networks, many new and important markets for this technology have been created, particularly in multinational corporations and multilingual nations. For those people whose job it is to follow world events as they happen and track the local response, there are strong incentives to use information filters that do not require full translation to access breaking news reports.

CLTR adds new twists to many traditional problems in information retrieval (IR) and computational linguistics. A CLTR system matches queries and documents in different languages by using bilingual dictionaries and thesauri for query translation or applying corpus-based techniques such as latent semantic indexing or cross-language query expansion. Dictionary-based systems must deal with translation ambiguity, and we have seen work based on word-sense disambiguation, part-of-speech tagging, syntactic analysis, query expansion, and Boolean models to address this problem. All the important user interface issues in IR, such as the organization of search results and document presentation, are magnified in the multilingual domain, where one generally cannot assume that the user is fluent in the document language(s). Our symposium also brought in researchers from the speech-retrieval community because speech retrieval (from text-based queries) and CLTR share the common property that the query and the information source do not have directly comparable representations. We even heard about some recent experiments in cross-language speech retrieval, a task that truly stretches the limits of modern recognition and retrieval technology.

Douglas Oard
University of Maryland

David Hull
Rank Xerox Research Centre

Intelligent Integration and Use of Text, Image, Video, and Audio Corpora

The presentations made at this symposium dealt with vastly different environments, ranging from digital libraries and broadcast news archives to virtual reality and, of course, the web.

There were two types of text, image, video, and audio (TIVA) integration: (1) data from one modality are used to improve and enable processing in another modality (leverage) and (2) multiple modalities were integrated to improve and enable overall system functions (concurrent integration).

Individual participants discussed the integration of images and text for finding captions, image retrieval, music retrieval, news broadcast segmentation, gaze tracking and speech, the analysis of video and audio, information retrieval and speech recognition, user studies, and formal logic to integrate image and semantic text representations.

Approaches to integrating TIVA sources included the following: First, features in separate modalities can be combined into a uniform vector representation suitable for learning. Second, in partially synchronized data, Viterbi alignment enables identification of correspondences across modalities. Third, heuristics can allow approximate matching across modalities. Fourth, hand-coded rules allow association of multiple feature types. Fifth, decision trees can learn the influence of one modality's data on another. Frequently, particular approaches could be applied in totally disparate modalities.

The presented research results could not be compared meaningfully because they were obtained on relatively small, private corpora. All participants agreed that the research would benefit tremendously from corpora that included multiple modalities, for example, a standard, large, public-domain database of text with pictures or a database of labeled broadcast video, including speech, music, and other sounds.

Alexander Hauptmann
Michael Witbrock
Carnegie Mellon University

Natural Language Processing for the World Wide Web

This symposium brought together researchers in natural language processing (NLP) from both academia and industry, including service providers on the World Wide Web. The purpose of the symposium was to discuss whether and how NLP techniques based on the syntax and semantics of natural languages could improve services on the web. A wide variety of technical approaches and application systems were demonstrated during the symposium, many of which were shown live off the web. The response to the symposium in terms of both papers submitted and the number of registrations and the enthusiasm of the participants were overwhelming.

It was encouraging to see practical applications of NLP for improved precision in document retrieval, natural language interfaces for better document search, semantic indexing of web pages, summarization and information extraction, machine translation of foreign-language web pages, question answering on the web, speech interfaces to web services, and dynamic creation of web pages using natural language-generation techniques. In addition, a position paper on the nature of the web for linguistic processing was presented.

Most of the discussions during the symposium focused on identifying key problems in web services where the application of NLP would make a difference. Many real-life problems in providing web services were presented by the participants from web service organizations. A consensus was reached that even if prior experiments had failed to show the benefit of NLP for information retrieval, the web offers many other critical problems for which NLP techniques are highly promising. Among such problems are query elicitation from users, question answering, summarization, topic spotting, machine translation, and document generation.

Kavi Mahesh
New Mexico State University

The purpose of the symposium was to discuss whether and how NLP techniques based on the syntax and semantics of natural languages could improve services on the World Wide Web.

Ontological Engineering

Explicit ontologies that support knowledge reuse, system integration, interoperability, and related tasks are becoming increasingly important. As ontology use becomes widespread, we must come to a greater understanding of the engineering issues relating to these ontologies. This symposium was convened to discuss current practices in engineering ontologies rather than the foundational issues raised at previous workshops.

A secondary goal of the symposium was to foster a stronger sense of community by sharing and drawing from our concrete experiences in developing and using ontologies. In the past, achieving this goal has been hampered by the interdisciplinary nature of the field, with research arising from knowledge base sharing and reuse, natural language processing, database schemas, and various standards efforts. For this reason, the workshop was organized around group discussions based on major topics and challenges in ontological engineering. This format encouraged participants to exchange the sort of critical information based in the experiences that often do not find their way into more formal papers and presentations.

Groups met to discuss the uses of ontologies; the reuse of ontologies; the form and content of ontologies; methodologies for designing, building, and evaluating ontologies; techniques for structuring libraries of ontologies; techniques for merging and integrating ontologies; and

tools for supporting ontological engineering.

Major issues to arise from the workshop were the increasing importance of merging ontologies and the problem of translating between ontologies. Many ontologies have been designed, ranging from ontologies for specialized domains such as medicine, law, products, and processes to the upper-level ontologies used in natural language processing. The challenge for the community will be the integration of all these efforts in such a way that we can realize the original vision of reuse and sharing.

There were stark contrasts among the purposes that participants used ontologies for and the pressures that these purposes placed on ontology design. For example, some participants required full axiomatizations to support automated inference, whereas others required broad coverage to support natural language processing.

In spite of the differences in purpose among participants, there was a surprising degree of consensus on the key issues and the need to focus on ontological content.

Adam Farquhar
Stanford University

Michael Gruninger
University of Toronto

Qualitative Preferences in Deliberation and Practical Reasoning

This symposium, attended by an audience representing AI, economics, decision analysis, and philosophy,

helped to establish the viability of qualitative decision theory as an area of AI (and perhaps, eventually, of economics). The papers and discussion at the workshop showed enough promising ideas to sustain this emerging field for many years to come.

Of the three elements in belief-intention-desire agent architectures, desire is the least studied in AI. Relatively little work has been done on how desires are rationally maintained, how they affect reasoning in general, and how we can infer the desires of agents and reason about them.

Microeconomics provides a well-developed theory, essentially unchanged since its formulation by L. J. Savage in 1954, of rational preference based on Bayesian probability and numeric utilities. This theory presupposes that all values are completely commensurable. AI cannot presuppose this commensurability in acquiring and working with incompatible goals and incomplete preferences and has sought a qualitative theory of preference to facilitate reasoning about partial preferences and conflicting desires, facilitate extracting and learning human preferences, and help guide the design of automated decision support systems. The symposium included presentations on semantic and axiomatic foundations; preference elicitation; preference revision; relations between argumentation, negotiation, and preference; and applications in medical therapy and product selection.

The most common subtheme consisted of attempts, some exploiting the von Neumann-Morgenstern probability-utility duality, to adapt methods for representing probability, possibility, nonmonotonicity, and conditionals to represent preference and utility. For further information, see www.kr.org/qdt.

Jon Doyle
Massachusetts Institute of
Technology

Richmond H. Thomason
University of Pittsburgh