

Reports on the AAAI Fall Symposia

5–7 November 1999

The 1999 American Association for Artificial Intelligence Fall Symposium Series was held Friday through Sunday, 5–7 November 1999, at the Sea Crest Oceanfront Resort and Conference Center. The titles of the five symposia were

- Modal and Temporal Logics-Based Planning for Open Networked Multimedia Systems
- Narrative Intelligence
- Psychological Models of Communication in Collaborative Systems
- Question-Answering Systems
- Using Layout for the Generation, Understanding, or Retrieval of Documents

This article concludes with a previously unpublished report on the 1998 AAAI Fall Symposium on AI and Link Analysis.

Modal and Temporal Logics-Based Planning for Open Networked Multi- media Systems

This symposium provided a forum for researchers involved in using formal methods and in design of networked multimedia systems and adaptive-reactive systems to identify common ground, relevant experiences, applications, open problems, and possible future developments.

To support intelligent and interactive multimedia applications, there's a need to tailor systems to possess and use knowledge about the application domain, user-requirement tasks, the context of interaction, communication, and performance parameters. Specific applications addressed concern video on demand, multimedia

presentations, multimedia synchronization protocols, multimedia collaboration, electronic commerce, nomadic services, and so on. Temporal and modal logics have been used to reason about time, action, and adaptive change and to program and verify networked systems.

Several topics were presented and discussed: (1) modal temporal-logic-based management of open distributed-processing systems and active networks; (2) modal-temporal-logic

The 1999 American Association for Artificial Intelligence Fall Symposium Series was held Friday through Sunday, 5–7 November 1999, at the Sea Crest Oceanfront Resort and Conference Center. The titles of the five symposia were (1) Modal and Temporal Logics-Based Planning for Open Networked Multimedia Systems; (2) Narrative Intelligence; (3) Psychological Models of Communication in Collaborative Systems; (4) Question-Answering Systems; and (5) Using Layout for the Generation, Understanding, or Retrieval of Documents.

specifications of service quality in multimedia networks; (3) combination of temporal-modal logics with formal languages such as LOTOS and CCS; (4) design of modal-temporal logics: operational models and implementation techniques, programming support and environments, and comparative studies of languages; (5) specification and verification technology: languages and tool support, model checking, and so on; and (6) software processes based on formal methods: design, analysis, verification, refinement, and so on.

Fawzi Daoud
GMD

Narrative Intelligence

People are narrative animals. By telling stories, we make sense of the world. We order its events and find meaning in them by assimilating them to more or less familiar narratives. A rich research tradition in a variety of fields suggests that narrative is fundamental to human intelligence. At the Narrative Intelligence Symposium, over 50 researchers gathered to discuss systems at the confluence of AI and narrative, systems that support, model, or investigate human narrative intelligence.

A wide variety of systems were presented: story generation, interactive fiction (including the first public demonstration from Joseph Bates's company Zoesis), systems that support storytelling by children and adults, agents as characters in narratives, agents that can tell stories, and more. These systems were built by researchers working at the intersection of AI and a variety of humanistic disciplines, including literary and cultural theory, art, fiction writing, psychology, and learning theory. Narrative was understood by these researchers from a variety of perspectives, including narrative as story, folk tale, small talk, drama, and world view. Despite this heterogeneity of application areas and approaches, we quickly developed a sense of community and common purpose as we realized we were addressing a set of common questions:

How can we create characters from which interactive narrative emerges? How can we create systems that can understand and generate compelling narratives? How can we build up narrative from a collection of elements, whether language, images, or plot? What properties does a system need to trigger narrative interpretation in the user?

At the same time, as a group, we began discussing broader questions of the relations between our systems and culture or society as a whole: What happens when agents can use narrative in a manner indistinguishable from humans? Is this a good goal? How can systems help us learn about and communicate our values

To support intelligent and interactive multimedia applications, there's a need to tailor systems to possess and use knowledge about the application domain, user-requirement tasks, the context of interaction, communication, and performance parameters.

and identities? How do new interactive media resemble and differ from more traditional media? How is AI technology changing these media? What role can narrative-intelligence systems play in enhancing everyday life? How can we build systems that play a meaningful role in our society? How can funding be found, both from the technical community and from the arts and humanities, to support the development of narrative-intelligence systems.

By the end of the symposium, it was clear that a new community is forming in AI, a community deeply informed by traditions in the arts and humanities and anchored in the construction of AI systems.

Michael Mateas
Carnegie Mellon University

Phoebe Sengers
GMD

Psychological Models of Communication in Collaborative Systems

Recently, the fields of AI, human-computer interaction, computer-supported cooperative work, and computational linguistics have spilled much ink on the topic of collaborative systems. Collaborative systems act either as partners to human users or as mediators between communicating users. The inventors of collaborative systems, being human, all have intuitions about what collaboration is and how it should be supported; so, many collaborative systems embody their intuitions. However, there has been an increasing focus on data—from looking at corpora of language use or

behaviors during collaboration to using systems. This link between data and systems is bidirectional: The data inform the design of systems, and the systems yield more data that highlight basic findings (about cognition, interaction, or the human use of technology) or applied findings (such as whether the systems are easy to use or whether they do what they're supposed to do). The third part of the triangle is theory. In the field of psychology, researchers generate theories to explain interesting observations, and then they generate more observations to test these theories—another bidirectional relationship. What is/has been/should be the relationship between psychological theory and collaborative systems? This question was the inspiration for this symposium.

The symposium was attended by a diverse international group in which social scientists were well represented; about 40 percent of the 50 registered participants were card-carrying psychologists (by way of their training or departmental membership). Overwhelmingly, it was an interdisciplinary group: When asked to characterize themselves as building things or studying things, many characterized themselves as doing both. A substantial number of the presentations focused on H. H. Clark's grounding theory and related ideas. These included not only attempts to formally model or implement the parameters of grounding but also experiments looking at communication in pairs or larger groups, through different media, with different tasks, among participants in diverse roles, and so on. Clark himself argued that communication with virtual partners is another form of disembodied language use, not

unlike what happens when people read text or interpret language at a time when it is not being produced by an actual speaker. Finally, several systems were described and demonstrations were presented of people interacting with and through collaborative systems along with discussion of the psychological assumptions and implications of these systems.

The symposium focused mainly on a single, comprehensive theoretical framework, Clark's grounding model. Other interesting approaches were sampled briefly (for example, gesture and language as an integrated system, speech acts, ACT-R as a modeling framework, the impact of cue phrases on comprehension, approaches to tutoring and pedagogy, and general contributions from the field of conversation analysis).

Susan Brennan
SUNY at Stony Brook

Alain Giboin
Inria

David Traum
University of Maryland

Question-Answering Systems

This symposium brought together researchers in knowledge representation, databases, and natural language processing interested in various forms of question answering. There were three focus areas in the symposium: (1) knowledge bases, (2) question-answering techniques, and (3) evaluation of question-answering systems.

Making all human knowledge available to ordinary people was identified as the grand challenge in the opening and vision session of the symposium. The internet has clearly contributed to the availability of knowledge to masses, but effective and efficient retrieval of this knowledge still poses significant obstacles. The ability to view the text documents as a knowledge base and exploit the semantics in form of linguistic knowledge and common-sense knowledge is the avenue where the community present at the symposium can make contributions. The vision panelists argued that question answering should not be just one-way

communication. Instead, the system should ask questions back from a user and engage the user in a conversation in the process of answering a question.

Two distinct models of question answering were discussed at the symposium. Under the first model, the system produces a coherent, concise, and focused explanation of the question based on a detailed representation of the topic on which questions are to be answered. In the second model, the system does not (always) directly answer the question but points the user to a text document that might contain the answer. Directing the user might be done by processing a text document for simple fact extraction or processing multiple sentences or documents to produce an answer. Even though there is a potential for combining the two approaches, not much seems to have happened on this front to date.

Knowledge bases with substantial linguistic knowledge are now available and were discussed at the symposium. For example, WORDNET provides elaborate semantic organization of lexical concepts. It includes knowledge about synonyms, hypernyms, and so on. WORDNET has been successfully used in several projects. WORDNET can be further enriched, for example, by disambiguating definitional glosses. The conceptual taxonomy developed at Sun Laboratories, in contrast to WORDNET, is based on the premise that there are no true synonyms, and the fundamental relationship between word phrases is subsumption. This taxonomy encodes subsumption relationships among about 15,000 words that are used as a basis for inferring subsumption relationships for new word phrases. The conceptual taxonomy used in conjunction with morphological knowledge and information-retrieval methods improves retrieval performance by about 20 percent over conventional techniques.

Substantial progress in deductive inference technology for question answering has been made. For example, industrial-strength systems to reason with recursion and negation are now available. Such systems have been used in modeling system properties such as safety and liveness but are

... layout influences the reader's assumptions about the role of each part of the text, grouping items together; highlighting; or marking items with a special purpose, such as headings and captions. These crucial signals can be lost when a document is transformed from one medium to another ...

yet to find extensive use in knowledge-based question answering. Unconventional models for finding answers, such as hypothetical answers, are still at a research stage and are yet to be tested experimentally. Some of the linguistic knowledge is beginning to find its way into commercial products. For example, Microsoft's ENGLISH QUERY product allows a user to create an English representation of queries on SQL (standard query language) tables and uses linguistic knowledge to answer variations on these English queries.

Formally evaluating and defining the semantic complexity of question-answering systems remains a difficult challenge. Preliminary proposals to measure the semantic complexity of question-answering tasks were made at the symposium. The semantic complexity of a question is often a function of the closeness between the question and the representation of the relevant knowledge in the knowledge base. A seemingly complex question can be answered quite easily by explicitly representing the answer or templates of answers in the knowledge base. Thus, any measure for the complexity of a task should take into account both the question and the processing necessary to answer it.

The ability of a system to answer questions is also a measure of the competence of the knowledge and the associated knowledge-processing algorithms in a knowledge base. A valiant effort to use question answering as a technique to measure the competence of a knowledge base was made during the Defense Advanced Research Project Agency's High-Performance Knowledge Bases Project. The question-answering task, in addition to serving as a measure of compe-

tence of a knowledge base, gave an impetus for studying the economics of the knowledge base-development process itself.

Question answering has been used extensively in several government-sponsored evaluations of AI systems, for example, the air-travel information system (ATIS), TIPSTER, and TEXT-RETRIEVAL Conference. Such evaluations have pushed the technology in a direction that has created several practically useful systems. At times, such evaluations have also swayed the research away from addressing difficult research problems that require a long-term research effort. Striking a balance between the requirement to produce results that are useful in the short run and the investment of a portion of effort for long-term research efforts appears to be the recommended recourse.

Vinay Chaudri
SRI International

Richard Fikes
Stanford University

Using Layout for the Generation, Understanding, or Retrieval of Documents

The symposium brought together 30 people interested in layout from various points of view. Four main themes were addressed: (1) the significance of layout in written language, (2) the automatic generation of documents with layout, (3) automatic formatting using rhetorical markup, and (4) the use of layout for information extraction.

The opening talk was given by an invited speaker, the psychologist Pat Wright (Cardiff University), a leading

Link analysis helps human analysts make inferences about organizations and activities based on lower-level records of people, places, things, and events.

expert on the ways in which document design can affect readability. Through examples, Wright surveyed the significance of layout from the reader's point of view. She pointed out that even without reading the words, a reader can often identify the genre of a document from its graphic design. On taking a closer look, the reader might find visual aids to navigation that save time and effort; this factor might be crucial in determining whether the reader goes on to read sections of the document or gives up straight away. During detailed reading, layout influences the reader's assumptions about the role of each part of the text, grouping items together; highlighting; or marking items with a special purpose, such as headings and captions. These crucial signals can be lost when a document is transformed from one medium to another, for example, when a formatted article is transmitted as a text file.

Throughout the symposium, our discussion was favored by the presence of Rob Waller (Information Design Unit, United Kingdom), who combines experience in the art of document design with a long-standing theoretical interest. In his presentation, Waller illustrated some reasons why document design so often fails, including a lack of reader feedback, poor command of the toolbox available to designers, and genre dissonance because of insufficient cultural awareness. Both Waller and Wright stressed during discussion that layout conveys meanings beyond those directly related to the text, just as tone

of voice in conversation can augment—or even negate—the literal meaning of what is said.

Several speakers were interested in the relationship between layout and rhetorical organization. Christophe Luc and Marie-Paule Pery-Woodley (IRIT, France) analyzed an example in which a complex rhetorical structure is expressed partly through a layout device (enumerated list) and partly through linguistic “discourse markers” (words such as *since* and *however*). Judy Delin (University of Stirling) showed an example of rhetorical relations across modalities, analyzing in particular the communicative function of illustrations in a bird book; Nancy Green (University of North Carolina) also addressed cross-modal relationships in a system that generates text integrated with graphics (tables, bar charts, and so on).

Another intersection of interests was the question of markup: how to mark up texts for layout and how existing markup in HTML or XML could be exploited for automatic formatting or information extraction. James Curran (University of Sydney) described an impressive system that learns rules for deriving an XML content structure from a source marked up in HTML. Alexander Kroener (DFKI, Germany) and Isobel Cruz (Worcester Polytechnic Institute) addressed problems of automatic formatting. There was general interest in the project of building corpora marked up for layout: Nadjat Bouayad-Agha (University of Brighton) described some problems in marking up a corpus of patient information leaflets, and Matthew Hurst (University of Edinburgh) discussed the special difficulties posed by documents that include tables. Thomas Kieninger (DFKI GMBH, Germany) presented an interesting system that automatically assigns tabular structure to a source document scanned in through optical character recognition or, alternatively, to a source laid out in ASCII with no markup.

All presentations at the symposium, including the invited talk, will be preserved in a AAAI technical report to be published in 2000. It was also decided that a web site would be set up, giving contact details as well as links to pub-

lications and other resources concerning layout; our thanks are owed to John Willmore (BIZINT Solutions), who offered to host this site, and Hurst, who will maintain it.

Donia Scott and Richard Power
ITRI, University of Brighton

AAAI 1998 Fall Symposium Report: AI and Link Analysis

The Symposium on AI and Link Analysis examined the uses of AI technologies to enhance a common method for examining large numbers of interconnected heterogeneous records. Link analysis helps human analysts make inferences about organizations and activities based on lower-level records of people, places, things, and events. Current uses of link analysis discussed at the symposium included discovering potential collaborators by analyzing web-browsing patterns, interpreting textual documents, detecting illegal activities of terrorist groups and criminal organizations, analyzing transportation networks, and discovering novel relationships among known medical facts. The ability of link analysis to represent relationships and associations among objects of different types has proven crucial to assisting human investigators in comprehending complex webs of evidence and drawing conclusions that are not apparent from any single piece of information.

The symposium drew over 50 attendees from diverse communities, including researchers from university and commercial research laboratories; software developers; technical representatives from current providers of commercial link-analysis software and several consulting firms; and analysts from military, nonprofit, and commercial organizations. Fifteen papers were presented along with surveys on social network analysis, visualization techniques, and relevant AI technologies. All participants contributed to working sessions on current systems, AI technologies, applications, and research issues and future directions.

David Jensen
University of Massachusetts