Reports on the AAAI Spring Symposia

23–25 March 1998

The American Association for Artificial Intelligence, in cooperation with Stanford University's Department of Computer Science, held the 1998 Spring Symposium Series on 23 to 25 March at Stanford University. The topics of the eight symposia were (1) Applying Machine Learning to Discourse Processing, (2) Integrating Robotic Research: Taking the Next Leap, (3) Intelligent Environments, (4) Intelligent Text Summarization, (5) Interactive and Mixed-Initiative Decision-Theoretic Systems, (6) Multimodal Reasoning, (7) Prospects for a Common-Sense Theory of Causation, and (8) Satisficing Models.

Applying Machine Learning to Discourse Processing

The purpose of this symposium was to address how machine-learning techniques can be applied to problems in discourse processing, following recent successes in the application of such techniques in other areas of natural language processing (NLP). Most attendees were computational linguists who are among the first to apply machine learning to discourse problems. The program included presentations on applying learning to a variety of problems, for example, discourse cue interpretation and generation, dialogue act recognition, discourse recipe acquisition, discourse segmentation, noun-phrase and verb-phrase anaphora resolution, and discourse strategy effectiveness, using a number of methods, including explanation-based learning, transformation-based learning, decision trees, genetic programming, and hidden Markov models. It also included presentations on methodological issues such as data-annotation schemes and use of data-visualization tools on corpora. In addition, two tutorials provided an overview of various machine-learning techniques and how some have been applied to other areas of NLP.

During the discussion and panel sessions, a number of open problems were raised. There was much discussion on the availability of annotated corpora in the public domain to facilitate the application of supervised machine-learning techniques and allow the comparison of results obtained using different learning approaches. Other issues included what the potential is for (or the limits of) the usefulness of applying learning methods; how they can be integrated with deeper discourse processing models; and what error rates can be considered acceptable, in particular, when the component is embedded in a larger system. Although these discussions did not yield conclusive results, they brought to attention many problems that need to be addressed to advance research in this area.

The American Association for Artificial Intelligence, in cooperation with Stanford University's Department of Computer Science, held the 1998 Spring Symposium Series on 23 to 25 March at Stanford University. The topics of the eight symposia were (1) Applying Machine Learning to Discourse Processing, (2) Integrating Robotic Research: Taking the Next Leap, (3) Intelligent Environments, (4) Intelligent Text Summarization, (5) Interactive and Mixed-Initiative Decision-Theoretic Systems, (6) Multimodal Reasoning, (7) Prospects for a Common-Sense Theory of Causation, and (8) Satisficing Models.

Nancy Green
Carnegie Mellon University
Jennifer Chu-Carroll
Bell Laboratories, Lucent Tech.

Integrating Robotic Research: Taking the Next Leap

AI robotics research has reached the stage where researchers must begin to work together on integrating their systems. In this symposium, the organizers were not so much interested in "what" research is being done as in "how" it is being done. The goal was to generate interest in integrating solutions across research groups, moving work in narrowly focused areas into integrated research efforts.

The symposium was structured around several major areas, including formalisms, metrics, architectures, abstraction, simulation, learning, competitions, grand challenges, and funding mechanisms. In early position statements, several researchers examined why integrated AI robotics research has not been happening faster. They attributed the slow progress to such factors as early fragmentation of the community; lack of reproducibility of experiments; lack of adequate formalisms; and insufficient funds, time, and motivation. Later talks examined issues such as how learning can help to generalize robot processes and how hardware abstractions can lead to reusable components.

Progress was made on several fronts. Salvatore Desiano will begin compiling a lexicon, which will allow researchers to begin talking the same language (or at least be able to translate between each other's languages). Several researchers, including Tucker Balch and Don Brutzman, agreed to make their robotics simulators publicly available so researchers can experiment in common environments. Alan Schultz offered to maintain a repository of reusable robot code. Greg Dudek, a 1998 AAAI Mobile Robot Competition organizer, listened to participants offer ways to use the competition to

further integrate AI robotics research. Virtually all the participants agreed that as a community, we must work to modularize our code and carefully document the module interfaces. A full summary of positions and action items is available at the symposium web site: www.aic.nrl.navy.mil:80/~schultz/aaai98/.

David Kortenkamp
Metrica

Alan Schultz
NRL

Intelligent Environments

This symposium was the first of its kind to assemble an audience from the AI and human-computer interaction communities to address the creation of *intelligent environments*. These are highly interactive spaces that use embedded computation to observe and participate in the normal, everyday events occurring in the world around them.

The goal of this symposium was to bring together groups working on nascent intelligent environments in academic and corporate research laboratories. Creating an intelligent environment requires a unique breadth of knowledge that crosses the traditional boundaries in AI research areas. Many communities were therefore represented at this symposium, including computer vision, natural language processing, multimodal interfaces, machine learning, and AI architectures. We had a valuable and unique opportunity to exchange both theory and practice based on accumulated, albeit early, experience. There was also clear sentiment that a new research community was in the process of forming.

Many of the presentations focused on specific environments that had been constructed for exploration and experimentation. These environments included intelligent homes, offices, classrooms, and command posts, many of which are in daily use and represent extraordinary engineering and integration efforts. A range of higher-level issues was also discussed that included sensory perception, modal subsystems, software architectures, and long-term applications.

There were several recurrent subthemes. The foremost of these themes was the need for better theories of context and intention. Even shallow theories could throw light on sensory data—enhancing their interpretation—and could help reduce the application specificity in so many of the presented systems. The role of machine learning in current and future intelligent environments to allow user customization and enhance reliability also received much attention. For more information, see www. ai.mit. edu/IE.

Michael H. Coen
Massachusetts Institute of Technology

Intelligent Text Summarization

The Spring Symposium on Automated Text Summarization is the third in a series of recent meetings on a topic that has been a dream of natural language processing researchers for more than four decades.

Recent advances in natural language technology, and more specifically corpus-based methods, as well as the abundance of text available on the World Wide Web, have brought summarization back into the spotlight after a period of diminished interest. Following on meetings in Germany (1993) and Spain (1997), the symposium had a number of excellent papers that mostly focused on the automated extraction of the most important concepts and facts from a text, allowing researchers from around the world to learn from each other and define the agenda for summarization research for the years to come.

There were six paper sessions—(1) Discourse, (2) Systems, (3) Evaluation, (4) Cohesion, (5) Compaction, and (6) Approaches and Methods—and two panels—(1) Statistical versus Symbolic Methods and (2) Evaluation. Each paper session contained presentations and discussion sessions.

A large portion of the symposium was devoted to evaluation. How can one quantitatively measure the quality of a summary? After several papers and panel presentations on the topic, an exercise was held in which all attendees participated. To compare and evaluate various evaluation techniques, participants (separated into six groups) each performed two evaluations of different summaries of two texts. Some summaries were created by humans, some by computer. After the participant session, the overall tendencies and cross-correlations of the evaluation scores given by participants were analyzed, presented, and discussed. In total, three different evaluation techniques were tested: one based on information theory (the Shannon game), one on task performance (the question game), and one on information retrieval (the categorization game).

A web page containing an extended bibliography of research in text summarization was set up and will remain available at www.cs.columbia.edu/~radev/summarization/.

Dragomir R. Radev
Columbia University

Eduard Hovy
University of Southern California
Information Sciences Institute

Interactive and Mixed-Initiative Decision-Theoretic Systems

Decision-theoretic techniques were originally developed to help make careful choices in high-stakes situations such as making large business investments, planning military strategy, and choosing among medical treatment alternatives. Researchers in the uncertainty in AI community are asking whether these techniques can also be applied to help make high-quality decisions in more commonplace situations where the stakes are lower.

Decision theory is an attractive framework for building interactive systems for decision making or decision support, but a traditional decision-theoretic analysis requires both a probability model and a utility model, and it is typically time consuming and tedious to elicit either one. This overhead might not be justified by the importance of the problem being solved, especially if the elicitation cost cannot be amortized over many problem-solving episodes.

Often, a single problem can be solved without eliciting a complete

model in advance. For example, an automated travel agent would not need information about all a user's travel preferences to build a single itinerary. Because it is usually impossible to ascertain ahead of time exactly what preference information will and will not be relevant to solving a particular problem, there is a need to interleave the elicitation of preference information with the problem-solving process itself.

Fortunately, the richness of the decision-theoretic framework provides valuable flexibility in problem representation. Unimportant portions of a problem space can be represented using coarse preference information, or omitted altogether, and more important parts of the decision space can be represented more precisely. Symposium participants presented several techniques for representing and computing with partial or abstract models. Various techniques were presented for eliciting the decision model incrementally, in conjunction with the problem-solving process. Well-established techniques from decision analysis, including sensitivity analysis and value of information calculation, were also discussed in the context of incremental model elicitation. Finally, the importance of self-explanatory systems was emphasized because the user needs to understand the impact of his/her communicated preferences and their role in the problem-solving process.

Peter Haddawy
University of Wisconsin-Milwaukee
Steve Hanks
University of Washington

Multimodal Reasoning

This workshop, chaired by Eugene Freuder, examined a variety of reasoning methods, including case-based reasoning (CBR), model-based reasoning (MBR), rule-based reasoning, constraint-satisfaction reasoning (CSR), and machine learning, in a wide spectrum of task domains, including design, virtual assembly, molecular scene analysis, legal reasoning, online help desks, and classification. The multimodal approach leads to improved overall problem-

solving performance as well as enhanced understanding and practice of the individual reasoning modes.

The methodology most often combined with another is CBR, particularly the adaptive variety, in which retrieval of similar past problem-solving episodes is followed by adaptation of the most promising of them to create a solution for the new problem. In this workshop, CBR is used typically through its retrieval phase—only what could be called CBr (CB, little r) for case-based retrieval—to find an initial ball park solution that is then adapted, for example, using MBR or CSR. In essence, CBr is used to locate a potentially fruitful part of the solution space in which to initiate further problem solving. CBR provides a way to compensate for weaknesses in a model, such as exceptions. A few systems plow back information harvested from newly solved cases into their model to update and enhance it. From the CBR perspective, MBR and CSR provide deeper understanding of how to accomplish adaptation and raise the issue of how the needs of finding good initial cases might be reflected in the indexes and similarity metrics used in case retrieval and, indeed, the case base. The synergy exhibited between CBR and MBR or CSR is typical of multimodal reasoning.

Not all projects involved CBR of course. For example, geometric and conceptual reasonings were combined in a virtual assembly project, which also used an opportunistic control architecture to combine different modes of reasoning. For the most part, however, the control architectures exhibited were straight line: call one mode (for example, CBR), then another (for example, MBR). Future work in multimodal reasoning will no doubt more fully address control and architectural issues as well as details of individual reasoning modes (for example, knowledge representation).

Eugene Freuder
University of New Hampshire

Prospects for a Common-Sense Theory of Causation

One of the central questions addressed at this symposium was how best to represent causal laws. A number of alternatives were discussed, ranging from the inclusion of an explicit primitive causal connective to the use of inductive definitions. Specialized action languages, applications of the situation calculus, and new frameworks in which to integrate ideas from probability theory were explored. Generalizations of structure-equation models from economics, including the axiomatization of mechanism-based theories, were also presented.

Discussions on how to connect laws to particulars and how to capture varying degrees of causal connectedness centered on approaches based on either counterfactual reasoning or a branching time ontology. Counterfactual reasoning was argued to have an important use in supporting learning from experience. Connections between a semantics for causation based on counterfactuals and one based on notions of sufficient and necessary conditions were also presented as well as issues having to do with causal explanation.

Discussions of future work drew emphasis on the need for connecting logical formalizations with nonlogical theories and also the need for elaboration tolerance if theories are to scale up. The difficulty in setting time lines for work in this area was also raised; in fact, it was pointed out that perhaps there was no single theory of causation but rather many smaller, local theories to be joined in some way, perhaps by analogy. There was almost overwhelming consensus that work in this area should begin to encompass more complex benchmarks, microworlds, and systems.

Charlie Ortiz
Harvard University

Satisficing Models

To effectively accomplish their goals, agents need to model their environment and other agents with which they interact. Building detailed, accurate, and up-to-date models, however, is a time-consuming activity and can detract from the actual problem-solving activities of the agents. The symposium participants discussed issues related to building satisficing models that enable agents to reliably perform at an acceptable level of effectiveness. The use of satisficing models is proposed to be a viable approach to generating satisficing behavior rather than optimizing behavior as observed in naturally intelligent systems.

The symposium included three invited talks given by Stuart Russell, Shlomo Zilberstein, and Edmund Durfee. Among satisficing models, Russell highlighted the use of learning to choose the best from a class of models with bounded execution time. However, one can use anytime inference algorithms with arbitrarily complex models. Rational metareasoning is required to make justifiable cost-benefit trade-offs on building and using such models. Zilberstein's position on satisficing behavior was the necessity to make explicit or implicit choice to seek suboptimal solutions using approximate reasoning, approximate modeling, optimal metareasoning, or bounded optimal decision making. Durfee espoused his position about multiple agents coordinating by knowing just enough about other agents and using default reasoning to generate expectations about others' behaviors.

Break-out discussion groups focused on specific issues such as cost-benefit trade-offs of building agent models, bounded optimality of satisficing models, and the updating of models of learning agents. The papers presented during the symposium discussed multiple aspects of satisficing behavior: negotiation and contracting mechanisms for reallocation and scheduling of distributed tasks, agents learning in electronic markets, reciprocation with other agents by building models of their strategic behavior, the development of satisficing agent-control mechanisms, the devising of belief-revision schemes with genetic algorithms, the use of games of deterrence for modeling business processes, the learning of heuristics in game playing, coordination without explicit communication, and so on.

Sandip Sen
University of Tulsa