AI Magazine Volume 15 Number 3 (1994) (© AAAI)

AAAI 1994 Spring Symposium Series Reports

■ The American Association for Artificial Intelligence (AAAI) held its 1994 Spring Symposium Series on 19–23 March at Stanford University, Stanford, California. This article contains summaries of 10 of the 11 symposia that were conducted: Applications of Computer Vision in Medical Image Processing; AI in Medicine: Interpreting Clinical Data; Believable Agents; Computational Organization Design; Decision-Theoretic Planning; Detecting and Resolving Errors in Manufacturing Systems; Goal-Driven Learning; Intelligent Multimedia, Multimodal Systems; Software Agents; and Toward Physical Interaction and Manipulation. Proceedings of most of the symposia are available as technical reports from AAAI.

Applications of Computer Vision in Medical Image Processing

There is a growing community of computer vision researchers who are working on medical applications. This interdisciplinary activity is in part application driven and related to the widespread proliferation of high-resolution medical imagers. It is also the result of increased interest by the medical community in image-based methods, especially in surgical applications and the study of the anatomical correlations of diseases in living subjects. The applied mathematical traditions of computer vision and robotics have proven useful in exploiting the rich information that is latent in high-resolution medical imagery.

Our symposium, which turned into a small conference, was attended by an exciting international mix of surgeons, medical researchers, computer scientists, and engineers representing prominent medical and technical institutions. A common theme that emerged in the workshop was that much of the work in the community is collaborative in nature because many players are needed to carry off a successful application.

The event opened with a keynote talk by Nicholas Ayache of INRIA. He discussed research tracks in computer vision applied to three-dimensional medical images, highlighting some of the similarities and differences with mainstream computer vision. He pointed out that worldwide, about $8 billion each year is spent on medical image production, and about $350 million each year is spent on medical image processing, with both areas growing rapidly.

The topics of the paper sessions included medical image understanding, medical image registration, image-guided surgery, anatomic modeling, medical image processing and analysis, magnetic resonance imaging, postprocessing of magnetic resonance data, and medical visualization. The quality of the presentations was consistently high.

The symposium included a poster session that was held to be successful, both technically and socially, by presenters and attendees alike. Another highlight of the workshop was a tour of the Robotics Laboratory that is part of Stanford University's Computer Science Department as well as the Stanford robotic radiosurgery system, which is an exciting development in medical robotics.

The level of enthusiasm at the workshop was high, perhaps because the event provided an avenue where medical computer vision was the main topic. Future directions for research in medical computer vision were the topic of a panel discussion led by Eric Grimson of the Massachusetts Institute of Technology. A rousing interaction occurred in which there was general agreement that the field is exciting and growing. Consensus was reached that the level of interest was high enough to justify holding a similar event next spring in Nice, France (3–5 April), to be hosted by Ayache.

William Wells
Harvard Medical School and Brigham and Women's Hospital

AI in Medicine: Interpreting Clinical Data

The subject of the AI in Medicine symposium was interpreting clinical data. In particular, the focus was on the development and application of AI technologies to the problems of interpreting and monitoring clinical data. Topics of interest ranged from real-time monitoring system design to knowledge discovery and machine learning. We distributed two data sets to participants in an effort to (1) emphasize the particular difficulties of managing real data acquired in clinical settings over more academic research issues and (2) provide standard data sets for two common clinical situations to allow researchers both a basis for methodological comparison and ease of access to those who are not familiar with the medical domain.

One of these data sets was collected in an intensive care unit on 1 patient over a course of 12 hours. This dense data set contained 120 K of data on approximately 30 variables and featured many idiosyncrasies common to real-world data: missing information, erroneous data, inaccuracies in time stamping, and so on. The other, sparse data set contained information on blood glucose levels, insulin dosage, and lifestyle data on 70 diabetic patients collected over periods of a few weeks to several months.

To our delight, most participants chose to exercise their methods on the provided data. However, even more pleasant was the finding that a few researchers decided to launch new projects based on the problems identified in the distributed data sets. The dense data set was more popular, possibly because problems such as ventilator management and data reduction-abstraction have traditionally been popular subjects in AI in medicine research.

We organized the symposium around three central themes: (1) interpretation of sparse (clinical) data, (2) interpretation of dense (critical-care) data, and (3) other critical issues. Issues such as erroneous, missing, or out-of-synch data were more pronounced in the critical-care situation, whereas the issues of modeling and control were emphasized by those who chose to tackle the sparse data set. A few consensus opinions emerged: It is often possible to generate and parameterize empirical models using clinical data. As shown by a number of participants, such models perform adequately for prediction tasks and, thus, offer a potential solution for problems such as automated control of treatment. Hence, the lack of precise physiological models should not be considered a limiting factor in applying AI technology to clinical problems.

Real-world data pose many difficulties in acquisition, interpretation, and presentation. Furthermore, data reduction and abstraction are important issues in attempting to limit information overload in clinical and critical care. Consequently, in this symposium, we saw progress in trend detection, temporal abstraction, selective monitoring, feature extraction, and data visualization.

Establishing and maintaining a patient context is essential for both interpretation and control tasks for both varieties of data. Many participants expressed difficulties in using the provided data when supplemental information was required to establish the proper context. Thus there is a trend in AI in medicine toward methods that modify their behavior based on contextual information.

The data sets distributed for the symposium, although far from being perfect, provided a valuable baseline for communication of research results among participants. This observation supports our belief that there is a need for well-annotated standard data sets for testing various AI approaches to the interpretation of clinical data.

We feel that the symposium reached its goals in identifying significant problems and bringing a diverse group of researchers together. For future attempts based on this model, we would advise more complete annotation of data sets and the definition of comparison metrics to help identify the strengths and weaknesses of various approaches.

Serdar Uckun
Stanford University

Isaac Kohane
Harvard Medical School

Believable Agents

The Believable Agents symposium was organized to help promote and focus the community's study of interactive believable characters. A *believable character,* a notion drawn from the arts, is one that seems sufficiently real and engaging that the audience can suspend its disbelief and react to the character in a direct and emotional way.

Building interactive versions of believable characters is an application of AI that is receiving increased attention from the entertainment and user interface communities. The goal is to create autonomous, interactive creatures that carry the same qualities as the noninteractive believable characters of traditional media. They must project a sense of being really there—aware, intentioned, rich in personality, and capable of significant social interaction.

The task of creating believable agents shares goals and methods with other AI research on autonomous agents. Relevant work occurs under the themes of situated agents and integrated architectures and across varied subareas, such as intention, emotion, and discourse. However, unlike much of this work, the goal here is not to build accurate psychological models of human performance or even to produce particularly intelligent or competent agents. The fundamental requirement is to achieve a persistent appearance of awareness, intention, and social interaction in an engaging, personality-rich creature.

One of the themes of the symposium was that AI researchers have worked to produce agents capable of reasoning, problem solving, learning by concept formation, and other qualities apparently central to intelligence. However, it is not clear how important these abilities are to creating believable interactive characters. Indeed, artists seeking to capture the essence of humanity in traditional believable characters have emphasized many other qualities, such as the expression of emotion and the quirks of personality. About half of the symposium was devoted to hearing artists speak about what makes characters real and considering how the goals for believable agent research might be different from the goals for other agent research.

Another theme was the distinction between copying reality and producing a subjective sense of realism. In the arts, this distinction is clear. For instance, impressionist paintings are able to convey an exceptional sense of realism without trying to be photorealistic. Producing believable agents seems sometimes to require realistic modeling of human minds and bodies and sometimes to depend on careful selection and abstraction of human traits. The judgments on how best to make these decisions seem to be artistic rather than technical.

The last topic discussed was how to measure believability and progress toward believability. Some participants wanted objective tests, but others sharply disagreed, noting that artists seem to have little use for scientifically valid techniques to evaluate the believability of their characters. The final view presented was that sufficiently new research topics demand the simultaneous development of new methodologies and evaluation criteria and that we might expect the same here given the depth of the merging of art and science.

Joseph Bates
Carnegie Mellon University

Computational Organization Design

AI has found numerous applications in supporting decision making in organizations but few in managing the complex issues of organizational design, analysis, reconfiguration, reengineering, and process change. Modern private and public organizations are facing immense pressures to rapidly reconfigure their processes, products, and relationships with other organizations. The cross-functional complexity of these changes—including their impacts on the technologies that organizations use; the structures of organizations; and the integration of human and cognitive issues such as skill requirements, cognitive loads, and performance-management systems—is immense. The problem of capturing and managing this complexity provides a tremendous opportunity for computational design-analysis support much in the way that other large, complex design and analysis problems (for example, architectural design, engineering design) have been supported by automated assistance.

At the same time, organization theories and design approaches have reached a degree of maturity that they can profitably be brought together in the new enterprise of supporting the reconfiguration of organizations. This new avenue, which we call *computational organization design* (COD), encompasses the theoretical, practical, and methodological aspects of AI, design, and organization theory.

The symposium was structured around three panel discussions, one each day. This structure was meant to reflect a representation and search orientation to COD, covering conceptual models and ontologies, COD design decisions and design knowledge, and architectures for COD design problems and systems.

Mark Fox, University of Toronto, moderated the first panel. Every design problem has to be represented, and a key question for COD is what organizational concepts and representations are appropriate for what problems and situations. The panel addressed alternative ontologies and ontological-conceptual modeling choices for COD. The validity of ontologies was discussed.

Les Gasser, University of Southern California, moderated the panel on design decisions and design knowledge. From an AI perspective, COD includes a conceptual model of organization that specifies a design space and a set of design decisions that refine and restrict the space. These decisions need to be informed by *design knowledge*, that is, knowledge about how to make choices among design space alternatives. One question the panel addressed was, What kind of specific heuristic, algorithmic, and design process knowledge do we have for different organizational domains and problems?

The final panel focused on the fact that AI-based design systems or design support systems need architectures to capture and manipulate conceptual models and design knowledge to make or support decisions. What are the practical, interesting, and relevant architectures (for example, constraint, search, dynamic simulation, and rule architectures)?

There were also seven sessions of presentations by the attendees, which generally followed the broad topics of the panels. Included were sessions on models of software organizations, simulation of organizations, and constraint-based organizational design.

For us, the most rewarding experience of the symposium was the recognition that many organization scientists are now doing sophisticated AI, and AI researchers are developing a deep understanding and practice of organization science. The degree of experience and comprehension across these several communities was remarkable and led to highly stimulating, substantive, and well-integrated discussions and debates.

Ingemar Hulthage & Les Gasser
University of Southern California

Decision-Theoretic Planning

Over the past few years, an increasing number of AI researchers have been applying decision-theoretic techniques to tasks that have conventionally been considered the province of AI planning. The time seemed ripe to convene a symposium devoted to this topic to share results, identify common technical issues, discuss potential applications, and attempt to establish a consensus on the definition and scope of decision-theoretic planning.

The symposium was attended by some 75 researchers from academia and industry, almost all of whom also provided material for the working notes. It was organized into six panel discussions and two poster sessions: The poster sessions allowed participants to present and discuss technical results, and the panel discussions confronted broader issues related to defining the discipline of decision-theoretic planning, identifying key technological issues, and matching technology with application domains.

Panels were held on the following topics: (1) a definition of decision-theoretic planning, (2) applications of decision-theoretic planning technology, (3) models of action, (4) models of preference and utility, (5) the application of classical planning techniques, (6) the use of abstraction to aid the planning process, and (7) decision-theoretic control of reasoning.

As the topics imply, most of the sessions were devoted to defining the problem itself and exploring the main technical issues confronting a researcher in building a decision-theoretic planning agent. One exception was the panel on potential applications of the technology. Areas of application included prehospital care of trauma patients, autonomous submarines used in data gathering and repair tasks, fault diagnosis, automatic scheduling of tasks to telescopes, organizational decision making, and oil-spill management.

Most of the discussion focused on how to incorporate richer models of uncertainty and preference (utility) into symbolic planning algorithms. One common point of debate involved the different assumptions researchers make about what information about the world the agent would get at execution time. At one extreme, a *plan* (a sequence of actions to be executed unconditionally) assumes that no additional information will be available; therefore, the agent might as well plan completely ahead of time.

At the other extreme, some researchers use fully observable Markov processes to model the agent, and instead of generating a sequence of instructions for the agent to execute, the output is a policy telling the agent what action to execute next depending on what state the world is in (this approach is also taken by those involved in reactive architectures). This approach assumes that the agent will automatically be given complete information about the world at every stage of execution.

Obviously neither of these extreme positions is realistic for most domains, but the discussion highlighted the importance of reasoning about execution-time information in generating plans (or policies). A unification of these two approaches might involve extending classical planning algorithms to reason about the agent's state of information, model actions that change its state of information, and generate plans that depend on information gathered at execution time. Likewise, the assumption of a fully observable world can be relaxed to produce execution policies that do not rely on perfect and complete information being provided at execution time.

In summary, the symposium's goals were well met: We have a much better overview of the field in terms of how various researchers define the problem of decision-theoretic planning, how they approach a solution, and what the potential uses for the technology are.

Steve Hanks
University of Washington

Detecting and Resolving Errors in Manufacturing Systems

Any system designed to perform a manufacturing task must have ways of detecting and recovering from errors. Timely detection of anomalies in the behavior of the system is essential for its continuous safe operation. Such detection involves preventing or minimizing the occurrence of faults through robust design, detecting abnormal conditions, isolating faults, and finding ways of maintaining safe operation despite the presence of faults. The manufacturing environment is often sufficiently uncertain and dynamic, and the manufacturing systems are sufficiently complex to make detection and recovery from errors a major task. This symposium was devoted to analyzing these issues and proposing how to create manufacturing systems that can achieve their tasks despite unpredicted contingencies.

The symposium brought together an international group of researchers having diverse backgrounds but sharing the common interest of applying AI techniques to the design and control of manufacturing systems. The openness of the participants and the congenial atmosphere allowed plenty of interaction and sharing of experiences and ideas that would foster continued collaboration. The emphasis was more on addressing real problems encountered with real systems and finding the right tool for the problem more than on proposing new unproven techniques. Panel discussions addressed techniques for monitoring and diagnosis, lessons learned from designing complex systems, and issues in executing schedules to achieve the desired plant behavior.

Considerable time was spent analyzing what features of the problem are relevant for the design of a robust system. Among the features discussed are the number of sensors available, the size of the plant, the number of states, the cost of the product manufactured, the cost of down time, the cost of failure, the resources available, and the control actions available. There is a trend toward performing more quality control during manufacturing, which requires the abilities to monitor processes and perform diagnosis in real time. A variety of diverse applications were mentioned, including manufacturing of appliances, monitoring of spacecrafts, monitoring of tool wear in milling machines, diagnosis of computer boards, handling of food by robots, monitoring of an automated bottling line, and scheduling of a hot steel mill.

Most of the participants advocated the use of models for monitoring and diagnosis, and through the symposium, a number of models were mentioned and discussed, including discrete-event systems, general diagnostic engines, Petri nets, neural networks, fuzzy rules, stochastic dynamic programming, and qualitative reasoning. A few people proposed more empirical approaches or almost purely statistical methods.

The ability to learn from experience or, at least, the ability to adapt to the environment or change strategies is relevant for most applications and was mentioned often. Substantial progress was achieved in monitoring and diagnosis, but recovery remains an open question, and the development of a general recovery methodology is still open. The perceived complexity of sophisticated modeling techniques is a major stumbling block that prevents user acceptance, and we need to find better ways of interacting with users. In conclusion, everyone agreed that in real applications, there is a continuous trade-off between flexibility and optimality, and cost is often a sobering concern.

Maria Gini
University of Minnesota

Goal-Driven Learning

Goal-driven learning refers to the process of using the overall goals of an intelligent system to make decisions about what should be learned, when learning should occur, and which learning strategies are appropriate in a given context. This focusing process can take place at any decision point during learning, for example, when formulating learning goals, scheduling learning tasks, selecting learning algorithms, pruning the space of theories to be considered, selecting a learning bias, or generating experiments for data gathering. Research in AI, psychology, and education has shown the need for intelligent systems to make decisions about what and how to learn. The common rationale and the principle around which the symposium was organized is that the value of learning depends on how well it satisfies the goals of the system.

The symposium brought together researchers from diverse research areas to discuss issues in how learning goals arise, how they affect learner decisions of when and what to learn, and how they guide the learning process. Neat utility theoreticians mingled with scruffy cognitive scientists and empirically minded machine-learning researchers to address common issues of goal-driven learning in a wide variety of task domains, including category formation, education, explanation, knowledge acquisition, manufacturing, problem solving, robotics, route planning, scientific discovery, and user modeling.

At the center of the productive dialog that ensued was a fundamental set of issues common to all: the nature of goals and their influence on the learning process. Much of the presented research was motivated by computational arguments: In any realistic task domain, time and resource constraints prohibit consideration of all but a few of the possible inferential paths. Consequently, any reasoner, human or machine, must focus attention and resources on pursuing those inferential paths that are likely to be most useful and, similarly, those types of knowledge acquisition, reorganization, or reformulation activities that are likely to result in useful learning.

Researchers discussed a range of goals, including task goals and learning goals, and goal-like influences, including biases and policies, and a corresponding range of computational models of goal-driven learning. The breadth of approaches and methodologies converged surprisingly well, demonstrating the theoretical generality of goal-driven learning and underscoring the timeliness of this symposium. A common theme was a model of planful learning in which systems actively decide what, when, how, and even whether to learn in the context of the overall performance task. The strength of such an approach to learning was demonstrated through theoretical arguments and empirical results from implemented learning systems. Researchers discussed broad frameworks for goal-driven learning as well as specific techniques and results to fill out the framework.

The symposium was an unabashed success: It is clear that goal-driven learning is an important recent advance in learning research with theoretical, as well as practical, implications—in technical terms, a hot topic. The symposium went a long way toward helping to define this new field and relate it to existing research in AI, machine learning, and cognitive science.

Ashwin Ram
Georgia Institute of Technology

Marie desJardins
SRI International

Intelligent Multimedia, Multimodal Systems

This symposium brought together an international group of researchers to discuss theoretical, architectural, and application issues in intelligent multimedia, multimodal systems. The speakers and attendees included researchers and developers from academia and industry and from a diverse range of academic and technical backgrounds. The working notes comprised 16 papers.

The symposium commenced with six papers on theoretical aspects of intelligent multimedia, multimodal systems. The session included discussions on the use of models and theories of natural language for interpreting and modeling multimedia, multimodal interaction. Issues raised also included the relevance of current models of discourse to structuring multimedia, multimodal interactions. The ALFRESCO system was described by Massimo Zancanaro to illustrate the relationship between natural language dialogues and hypertext navigation of a multimedia information space. David Koons discussed the use of gestures, gaze, and natural language in a simple spatial manipulation environment, making use of Jackendoff's ideas on conceptual semantics to provide a common representation across the different modalities. Elisabeth Andre and Thomas Rist discussed the use of rhetorical structures to plan and generate multimedia presentations. In the only empirical study in the symposium, Sharon Oviatt investigated the use of speaking and writing using a simulation environment to look at the strengths and weaknesses of possible future types of interactive system. A theory of modality was proposed by Niels Ole Bernsen that attempted to map task-domain information onto particular forms of multimodal interaction.

In the applications session, John Meech considered how models of users and tasks could be used to integrate multimedia and multimodal interaction in real-time systems. Similarly, Fergal McCaffery described how voice, graphics, and text could be added to an existing circuit-testing application using models of the user and the application. In a description of CHATTER, Eric Ly and Chris Schmandt explained how speech-only input could be used to access computer-supported work group information over the telephone.

A multimedia knowledge delivery system that used a strong representation paradigm was discussed by James Lester and Bruce Porter. This system provided multimedia explanations in the domain of biology and was based on single, rather than multiple, representations of the knowledge in the system. A military application in which natural language had been added to an existing graphics-based system was described by Elaine Marsh, Ken Wauchope, and John Gurney. The system (a tactical battle management system for the United States Navy) makes use of natural language discourse models to relate speech and direct-manipulation interaction.

In the session on architecture, Claudie Faure and Luc Julia described an agent-based architecture for multimodal interaction. The architecture allowed pen-and-speech–based input to be integrated to generate graphic drawings. Manuel Perez and Robert Jacob described the architecture used in the battle management system discussed previously. This architecture included a discourse module using a focus stack to relate dialogue focus during a multimodal dialogue.

Srdjan Kovacevic described how a model-based approach to multimedia interaction had been developed in TACTICS using blackboards for dialogue control and integration of modalities. Keith Werkman provided a proposed architecture for applying concepts from distributed AI to multimedia, multimodal systems. This architecture was closely related to the work of Dick Bulterman, who described the architecture used to implement and control presentation in multimedia information systems. His architecture included a constraint mechanism with both static and dynamic constraints. The architecture supported both authoring and running of multimedia applications.

The symposium was extremely enjoyable with high participation and involvement from all who were there.

Peter Johnson
University of London

Software Agents

Software agents are sensor-effector systems that operate within real-world software environments such as operating systems, databases, or computer networks. Their sensors observe features of this external environment, and their effectors can both alter the state of the environment directly and communicate with other agents. Software agents pursue goals such as acquiring information about the environment or modifying its state, either individually or in teams. In contrast to work on human-computer collaboration, our focus is on agents with a high degree of autonomy and flexibility.

Advances in computer, information, and telecommunications technology have made software agents both necessary and possible. However, the role of AI and AI researchers in these developments has yet to be determined. We believe that AI has the potential to contribute to these developments and that software domains offer fascinating research challenges. To capitalize on this opportunity, the Software Agents symposium brought together researchers in this new area to develop a common vocabulary and identify the fundamental research issues that define it.

The symposium participants discussed many questions: What exactly is a software agent? What are challenge problems (or "killer aps") for the community (for example, information retrieval from the worldwide web)? Will research results obtained in software domains generalize to physical domains? What possibilities exist for software agent interaction (for example, can the Bell Labs visitor-bot query the University of Washington's Internet softbot for information about visitors, for example, their telephone numbers)?

The symposium was attended by over 70 participants from major university projects, including Carnegie Mellon University, Stanford University, the Massachusetts Institute of Technology, and the University of Washington, as well as major companies, including Microsoft, AT&T, Apple, and IBM.

Oren Etzioni
University of Washington

Toward Physical Interaction and Manipulation

The range and scope of practical robotics applications critically depend on the ability of robots to physically interact with their environments. Current successful systems are specialized, and they typically involve carefully controlled, well-understood work spaces with little or no sensory feedback. Construction costs and inflexibility limit the reusability of these systems. The general manipulation skills of humans and other animals contrast starkly with the current capabilities of robots. Unlike most current robots, humans rely on rich sources of sensory feedback to cope with uncertainties in the world. Robotics researchers with a broad range of interests and experience discussed organizational principles for robust sensorimotor control in uncertain environments.

From the beginning of the meeting, the perennial debate questioned whether robotic agility will be achieved by general-purpose robots using complex flexible manipulators, for example, anthropomorphic hands, or by a suite of specialized robots equipped with interchangeably specialized manipulators. Consider the observation that $700 can buy a food processor, a mixer, a bread machine, a coffee maker, a juicer, a telephone answering machine, and a vacuum cleaner (at discount prices). Each device automates some portion of a task, limiting the interaction with the world to well-defined actions by requiring materials to be loaded and unloaded or the business end to be guided. The same need for quickly reconfiguring material-handling and material guidance systems is also critical for office autonomy and flexible manufacturing.

A couple of live robot demonstrations clearly drove home important points. Larry Leifer and H. F. Machiel Van der Loos opened their rehabilitative desktop assistant robotics laboratory to our symposium. An important lesson is that devices are more readily accepted if users can anticipate their behavior. Users of the desktop assistant were disturbed by complex robot arm motions that are required by joint singularities. Ken Salisbury brought his PHANTOM haptic feedback device to the meeting. Its high fidelity in expressing to the touch a simulated world clearly illustrated his observation that current robot actuators have orders of magnitude less dynamic range than biological muscles. A robot can lift 20-kilogram payloads, or it can apply gentle forces, but it cannot yet do both.

Building physical systems verifies models and assumptions, but physical competence requires a sound foundation of sensorimotor control. The challenge is to build general reusable components and combine them flexibly. A paradigm is needed for building new combinations of concurrent interacting subsystems and coordinating them so that they cooperate, rather than interfere, with one another. Discrete-event systems might be a good candidate for this role. Although complete general vision systems are still beyond the horizon, some basic visual abilities, such as motion and stereo vision, have been used in successful demonstrations of specialized tasks, such as grasping a ball floating in near-zero gravity. Competence in many unconstrained environments will require modeling. For example, compliant catching will require modeling motion and inertia because events occur too fast for feedback reactions alone.

David Coombs
National Institute of Standards and Technology

Steven Whitehead
GTE Laboratories