AI Magazine Volume 11 Number 3 (1990) (© AAAI)

AAAI 1990 Spring Symposium Series Reports

■ *The American Association for Artificial Intelligence held its 1990 Spring Symposium Series on March 27–29 at Stanford University, Stanford, California. This article contains a short summary of seven of the nine symposia that were conducted.*

The American Association for Artificial Intelligence held its 1990 Spring Symposium Series on March 27–29 at Stanford University, Stanford, California. This article contains a short summary of seven of the nine symposia that were conducted (reports were unavailable for the symposia Planning in Uncertain, Unpredictable, and Changing Environments and The Theory and Application of Minimal-Length Encoding).

AI and Molecular Biology

Although molecular biology was the domain of several early AI systems, most notably Molgen and Dendral, this area has been relatively quiet until recently. The last few years have seen an explosion of data, knowledge, and analytic techniques in molecular biology, which has triggered a renaissance of AI research in this domain. This growing community of computer scientists and biologists gathered together for the first time at the 1990 AAAI Spring Symposia. This symposium received applications from more than 150 researchers from nearly a dozen countries, representing ongoing research efforts at more than 30 institutions, including several government laboratories and industrial concerns.

The work presented spanned a wide range of topics, from basic research in machine learning and the automatic generation of representations to the application challenges of squeezing out that last 1 percent of error in automated interpretation of sequence gels. The AI techniques used by the presenters spanned an even greater range. They included A*, inductive category formation, Bayesian inference, computational linguistics (both applied to texts and to DNA sequences), neural networks, qualitative modeling, hierarchical pattern recognition, case-based reasoning, deductive inference in Prolog, Connection Machine-ism, model-directed visual processing, constraint propagation, expert systems, knowledge base maintenance, minimal length encoding, object-oriented databases, simulation, induction of context-free grammars, and various knowledge-acquisition technologies. I believe that more different AI technologies have been applied to molecular biology problems in the last few years than have ever been brought to bear in any other specific application domain.

This symposium not only covered the science currently being done but also what might be termed the sociology of the new community. Representatives of the National Institute for the Humanities, the National Science Foundation, and the Human Genome Project spoke about the kinds of grant money, fellowships, and training support available for AI and molecular biology work. In addition, four of the most senior members of this young field—David Searls, Doug Brutlag, Peter Friedland, and Joshua Lederberg—conducted a panel discussion about where the field has been and where it is headed. Their conclusions were widely shared, and the excitement evident at this meeting presages a bright future for the field.

—*Lawrence Hunter*
National Library of Medicine

Artificial Intelligence in Medicine

The AI in medicine (AIM) symposium was attended by researchers from the United States, Canada, Denmark, the Netherlands, Italy, and the United Kingdom. It comprised seven sessions: diagnostic systems, uncertainty management, intensive care unit (ICU) applications and monitoring, temporal and spatial reasoning, knowledge representation and modeling, knowledge acquisition, and evaluation. Because more than 60 posters were presented, we can only give a flavor of the issues that were raised.

In the session on diagnostic systems, several participants presented work on methods for handling multiple-disease diagnoses. The approaches were diverse and included methods that use abstraction and conflict resolution. Todd Johnson and Jack Smith (Ohio State University) presented a diagnostic architecture that applies a spectrum of abductive strategies for medical problem solving, offering a new approach to overcoming system brittleness. However, Mike O'Neil, Andrzej Glowinski, and John Fox (Imperial Cancer Research Fund, London) argued for a knowledge-intensive approach to overcoming brittleness in the context of their architecture for a decision-support system for primary care.

The majority of the work in the uncertainty management session dealt with Bayesian belief networks. Several researchers described new methods for the acquisition of belief networks. One such method (Stig Andersen, Aalborg University, Denmark) allows the construction of large belief networks from smaller belief network modules; another approach (David Heckerman, Stanford University) provides a local method for experts to incrementally specify a global belief network structure.

Numerous research applications of AI are in the ICU and patient-monitoring domains. Reid Rubsamen (Massachusetts Institute of Technology [MIT]) discussed building smarter patient monitors by extending the semantic descriptions of traditional hemodynamic waveform data to capture subtleties in the waveforms. Larry Widman (University of Texas, San Antonio) described work on a system for monitoring cardiac arrhythmias that applies knowledge to recognize and resolve artifactual monitor signals.

Temporal reasoning continues to be an area of particular challenge in AIM research. Luca Console and Pietro Torasso (University of Turin) described a technique that involves integrating causal models with temporal-constraint satisfaction to perform abductive problem solving. Alex Yeh (MIT) introduced an innovative qualitative method for finding the average rates of change in repetitive measurements, such as recordings of cardiac output.

Reasoning about three-dimensional relationships is central to both diagnosis and therapy planning. Ioannis Kapouleas (Rutgers University) discussed new methods for reconstructing three-dimensional structure from two-dimensional computed tomographic images. Jim Brinkley (University of Washington) described preliminary work on the development of a large distributed knowledge base of structural biology that could serve as a resource for the next generation of spatial-reasoning systems. Ira Kalet and Witold Paluszynski (University of Washington) described the Radek system, in which symbolic techniques are used to evaluate the merits of putative treatment plans. Jeffrey Berger and Kristian Hammond (University of Chicago) presented Roentgen, in which the enormous search required to construct an initial radiotherapy plan is neatly finessed through the use of case-based reasoning.

Research on the application of AI techniques to medicine continues to spark investigation into the fundamental principles of knowledge representation and inference. Thierry Barsalou (IBM) presented his research on the Penguin expert database system. In Penguin, knowledge is permanently stored in a relational database and is automatically mapped into an object-based representation at run time to support efficient inference.

Considerable attention was given to the relationship between learning by machines and learning by humans. Both Mitch Medow (University of Wisconsin) and Yeona Jang (MIT) presented learning architectures that model aspects of human learning. Medow's ambitious Hematmid system uses similarity-based learning to generate multidimensional, tangled concept hierarchies from the input of prototypical cases. Jang described an architecture for compiling causal knowledge into more efficient associational knowledge in the course of experience.

David Evans (Carnegie-Mellon University) raised the intriguing hypothesis that current trends toward the development of expert systems based on causal and probabilistic models will result in systems that are less psychologically real than those based on the empirical associations that seem to underlie most medical judgment. Evans suggested that because natural language processing requires psychological models of inference to establish contexts for discourse, many second-generation expert systems might be unsuitable for use with natural language front ends. Evans's Med-Chart project attempts to acquire the background knowledge required for interpreting the natural language of patient records.

Ted Shortliffe (Stanford University) moderated a panel discussion on the evaluation of medical expert systems. Evaluation studies of three large expert systems were presented. Ruth de Bliek (University of North Carolina) described the formative evaluation of the advisories produced by the Mentor system at Stanford, which warns physicians of potential errors in the administration of certain drugs. Hein Moens and Larry Kingsland (National Library of Medicine) discussed reasons for the decreased diagnostic performance of the AI/Rheum system when the program was presented with 1570 cases from a Dutch rheumatology clinic. Johan van der Lei (Erasmus University, Rotterdam) described the evaluation of the Hypercritic system, which critiques physician's management of patients with hypertension.

One fundamental issue recurred throughout the symposium: Although it is clear that the AIM community has contributed substantively to AI, the contributions to medicine are far less palpable. The consensus was that the primary obstacles were not scientific but rather logistical and political. Physicians are typically unwilling to perform extra work to obtain decision support. Busy health-care providers can rarely take the time required to enter patient-specific data into a computer to obtain a consultation. The frustration for developers, of course, is that such data are often already available online from office information systems and computers used by clinical laboratories. Participants debated the role that the AIM community should play in facilitating computer networking among hospitals, clinical laboratories, and physicians' offices and promoting standards for data interchange. Finally, there was the common expectation that AI will ultimately demonstrate important benefits for health care, although the dissemination of AIM systems and the measurement of their success represent considerable challenges.

—*Gregory F. Cooper* and
Mark A. Musen
Stanford University

Automated Abduction

Explanations are important in many different aspects of intelligence, so it is no surprise that key tasks in many different areas of AI can be viewed in terms of abduction or inference to the best explanation. In expert systems, the best known abduction problem is diagnosis. In natural language comprehension, plan recognition can be viewed as an abduction problem involving the inference of goals from observed behavior. In qualitative physics, postdiction is an abduction problem involving explaining states of the physical world in terms of processes and causal laws. In machine learning, explanation-based learning strategies improve performance using processes that construct explanations. Explanations are important in so many areas of AI that the problem of finding efficient abduction methods for constructing and evaluating explanations has been labeled an AI-complete problem.

Abduction-related work has been done in different areas of AI for more than a decade, but until recently, researchers working in a given subfield often failed to recognize that they might benefit from work on abduction by people in other areas. This symposium was aimed at facilitating cross-fertilization in the hope of accelerating research advances in all subfields of AI concerned with explanations. Researchers with interests in business, natural language processing, diagnosis, qualitative physics, machine learning, and discovery gathered to discuss the role of abduction in their disciplines.

A number of domain-independent methods for constructing and evaluating explanations were discussed, including generalized set covering, methods based on assumption-based truth maintenance system technology, methods based on explanation patterns, cost-based abduction, coherence-based approaches, and probabilistic approaches. Several general theories of abduction were proposed giving axiomatic characterizations of aspects of abduction, semantic theories of abduction, and analyses of the complexity of abductive computations. It is not surprising that abduction, like many other AI-complete problems, is intractable in general, but interesting recent results were presented that exactly characterize when and why abduction is hard.

It was encouraging to see that several general formal theories of abduction have begun to develop and more encouraging to see that these theories were closely tied to the algorithms being used in applications. No psychological data about how people construct and evaluate explanations and few formal evaluations or comparisons of alternative approaches or systems were presented. However, informal discussions comparing different approaches, methods, and implementations (for example, Bayesian probabilistic reasoning versus connectionist networks) took place. In addition, the relationships between abduction and various other forms of inference were explored, including deduction, induction, and many other forms of plausible reasoning.

—Paul O'Rorke
University of California at Irvine

Case-Based Reasoning

Case-based reasoning (CBR) addresses problem solving and reasoning tasks in domains where experience is strong, but the domain model is either weak or weakly exploited. A case base contains specific solutions or specific precedents that operate to solve problems or support relevant chains of reasoning. When a new problem or scenario is encountered, instances from the case base are retrieved and adapted to address the situation at hand. Tasks that lend themselves to CBR are typically characterized by multiple solutions where trade-offs are more important than optimization, and the ability to compare different possibilities is more important than the notion of correctness. Medical diagnosis, legal reasoning, and problems in design are prototypical examples of CBR tasks as well as memory-oriented planning and reasoning from analogies.

Six standard problems must be addressed in designing any CBR system: (1) the representation of cases in memory, (2) the organization of cases in abstraction hierarchies, (3) memory access through indexing and retrieval algorithms, (4) conflict resolution across multiple cases, (5) the adaptation of retrieved cases to address the current problem, and (6) the growth of the case base in response to previously solved problems or supervised learning. In addition to these issues that are specific to CBR, there is also a lot of interest in hybrid CBR systems that mix CBR with rule-based reasoning, explanation-based learning, and other AI technologies.

The format for our symposium was based on 11 paper presentations followed by prepared commentaries and a brief period for questions and discussion. In addition, panel discussions addressed CBR and machine learning, CBR and planning, CBR and intelligent tutoring, and progress on a generic architecture for CBR. Finally, seven papers were published in the working notes but were not presented. We specifically targeted contributions to basic research rather than systems per se; so, our orientation was slanted toward issues and ideas. Most of the work reported focused on case representation, organization, and retrieval. To give a feel for what went on, I'll mention a few specific issues and contributions.

Effective case retrieval depends on careful indexing strategies. By now, everyone seems to agree that feature-based indexes are not adequate for retrieving structured cases. Unfortunately, pattern matching based on shared structural similarity is NP-complete, so we need to find some cleverness for dealing with structural matching. Karl Branting proposed an indexing scheme based on structural difference links in an effort to make retrieval tractable. Jeremiah Faries argued for a different tack by stressing the role of derivational analogies in political reasoning. We might not have a consensus on indexes, but we have many ideas that are operational over a variety of applications. We use feature-based indexes, structural indexes, and cross-context indexing. Roger Schank's group at Northwestern University introduced the universal index frame as a first step toward identifying a universal content theory for indexing.

The term content theory popped up frequently at this meeting, although I have not often seen the term in print. It does not have quite the same meaning as domain theory, but it is connected to Newell's idea of the knowledge level and what I would more specifically call domain-specific semantics. Somebody is going to have to nail this term in writing if it is going to be meaningfully used.

A number of people addressed the question of how to optimize case abstraction: We want to find a level of description that balances predictiveness and predictability. Doug Fisher explained how this problem is tackled in concept formation systems and extended his view of the problem to CBR systems. William Mark introduced an orthogonal issue and argued that it might at least be as important to worry about how cases can be broken into component pieces to provide an optimal granularity for the case base. It seems safe to say that we are still somewhat undecided on the general issue of what constitutes good case encoding and organization.

All in all, this gathering reflected the enthusiasm and momentum one hopes to see when a new technology starts to blossom. We regret that many publishable papers were rejected to keep the meeting small, but it is exciting to see how rapidly the CBR community is expanding. This symposium was a wonderful opportunity for us to have a small gathering before our numbers prohibit us from even trying.

—Wendy Lehnert
University of Massachusetts

Knowledge-Based Environments for Teaching and Learning

This symposium focused on the use of technology to facilitate learning, training, teaching, counseling, coaxing, and coaching. Sixty participants from academia and industry assessed progress made to date and speculated on new tools for building second-generation systems. The selection of topics and participants was motivated by a desire for ideological breadth and depth. Panel leaders included William J. Clancey and Alan Lesgold (real-world systems), Kurt VanLehn (cognitive models), Beverly Woolf (discourse systems), Elliot Soloway (alternative environments), and Sarah Douglas (supportive systems).

Researchers have moved away from building omniscient tutors capable of detecting all possible errors and misconceptions. Instead, research is now focused on building empathetic partners that choose from among several forms of interaction based on the content of the communication and the needs of the student. Possible communication styles include didactic explanation, guided discovery learning, coaching or coaxing, and critiquing. Although no one style is preferred, different tutorial applications will be better addressed with a given primary style. In addition, the style of interaction varies within a tutorial domain as well as across types of domains.

Pressing research issues in human-computer communication were identified in both AI and education. In AI, issues include the representation and control of knowledge, the choosing and organizing of domain knowledge, and the characterization of coherence in machine response. Educational research issues focus on adequately modeling the student and the pedagogical context and recognizing how a system might stimulate and facilitate the student's own abilities and creativity. A separate issue concerns how relevant knowledge should be presented once it has been selected.

The cognitive modeling group argued for increased use of modeling at three stages of knowledge-based system design, primarily (1) development of pedagogical and subject-matter theories, (2) design of instruction, and (3) delivery of instruction. Of these phases, the design of instruction is the one that seems to have achieved the most direct benefit from cognitive modeling, including substantial benefits from modeling subject matter experts. However, work on modeling good teachers and tutors has only just begun (with the exception of a few early classics). The actual delivery of the instruction is the area where cognitive modeling has found the least fruitful application.

VanLehn underscored the fact that modeling is just good engineering practice, regardless of whether one is building a hydroelectric dam or a science course. With tongue in cheek, he suggested that if students could sue malfeasant instructional developers, cognitive modeling would be much more common because it is so obviously effective. Clancey, however, was more reserved about the utility of cognitive modeling.

In the move away from building all-knowing and all-powerful tutors, researchers have focused on developing environments that implicitly elicit information about student goals and plans. Human dialogue succeeds despite ambiguity and digressions because both participants model the discourse, the subject matter, and the other speaker, and both actively work toward success of the discourse. This situation suggests that continuing efforts be made to enhance the machine's ability to do its part. Techniques such as plan recognition and learning still play only a small role in current teaching systems.

Clancey and Lesgold led several discussions on the impact of knowledge-based systems in industry and the military. The clear emergence of new architectures and positive training results has produced the feeling that progress is being made. Indeed, several systems were described that achieve the two-sigma effect, which is the same improvement in learning that results from one-on-one human tutoring over classroom tutoring. Several success stories were described in which students using tutors learned knowledge and skills in one-third to one-half the time it took for a control group to learn the same material.

Several areas emerged as hot or new research areas. Situated learning (and teaching-acting-planning) frequently arose as a topic. It was espoused primarily by Clancey, Jeremy Roschelle, and Etienne Wenger (Institute for Research on Learning, Palo Alto, California). Because situations or contexts in which a skill is learned cannot be exhaustively or completely described, training systems inevitably predetermine what is relevant. Similarly, conventional AI models of expertise omit how experts know what is relevant and how they change their minds. This approach suggests that AI systems need to place increasing emphasis on representation as an activity within a perceptual space and organized by social interaction. Current systems omit the social context in which domain representations are created, justified, and changed. Currently, knowledge-based cognitive modeling cannot characterize the work somebody must do to understand specific artifacts or tools of a community. Other important research areas were the computer as mediator, the empowering of curriculum designers, and qualitative reasoning.

Participants represented diverse backgrounds and methodologys: thus little commonality might have been expected. However, a small consensus was achieved and some new scientific ground broken. Agreement was reached on the need for a variety of discourse approaches and the need for cognitive models, although no single solution to achieve widespread use of either was forthcoming. In addition, basic research is needed in planning and plan recognition; the building of natural language interfaces; and the application of architectures, such as blackboards, to teaching systems. From the viewpoint of communication, the symposium was a real success; discussion was lively and at times controversial. Research appears to be strong in depth, broad in perspective, and motivated by the promise of building more powerful teaching environments with greater knowledge, increased inference capability, and more complex reasoning ability.

—*Beverly Woolf*
University of Massachusetts

Knowledge-Based Human-Computer Communication

This symposium addressed the important goal of how to understand, design, and build cooperative problem-solving systems, acknowledging that most knowledge-based systems are intended to assist human endeavors and are almost never intended to be autonomous agents. The symposium questioned the basic assumption that the most widely publicized goal of AI (understanding and building autonomous, intelligent, thinking machines) is also the most important one.

The discussion centered around eight major themes: (1) agents versus augmentation, (2) user models: what they are and why, (3) knowledge representation at the user interface, (4) the impact of this research in the real world, (5) presentations and representations, (6) dialogue issues in collaborative systems, (7) the sharing of the load between humans and computers, and (8) tutors versus critics. The symposium turned out to be an interesting and exciting forum for these issues and succeeded in building bridges between interdisciplinary research efforts.

—*Gerhard Fischer*
University of Colorado at Boulder

Text-Based Intelligent Systems

The motivating theme of this symposium was the rapid advance of the availability of online textual information as opposed to the extreme difficulty of developing large knowledge bases. This contrast suggests a need for AI systems that derive their power from large quantities of stored text rather than hand-crafted rules. Such text-based intelligent systems must combine AI techniques with more robust but shallower methods of processing and accessing texts.

This symposium brought together about 70 scientists, who focused mainly on the issues of text processing, information retrieval, and applications. Active discussion sessions followed short position papers, with discussion accounting for half of the meeting.

Several themes emerged as the focal points of research in text-based intelligent systems (TBIS). First, the attendees shared a recognition that the traditional, knowledge-intensive methods of AI in general and natural language processing in particular cannot apply to large volumes of real text. Many of the position papers thus tried to marry weak methods, such as statistics and word-based text analysis, with more knowledge-based techniques, such as syntactic parsing and semantic interpretation.

Another central theme was robustness. Much of the research addressed the design and implementation of robust text processing systems, although the definition of robustness ranged from the handling of arbitrary text within a constrained domain to millions of words of running text. There was also a fair amount of debate on the depth and accuracy required for a system to be considered robust.

Related to these two themes was the emerging importance of corpus-based analysis in natural language processing. By their nature, weak methods are empirical and must be developed with respect to some body of examples. Similarly, robust methods can only be evaluated on specific tasks on a particular set of texts, or corpus. Large corpora thus play a role in the design, development, and testing of text-based systems.

In addition to cross-disciplinary exchanges on some of the issues, the symposium served as a starting point for resource sharing relative to text-based systems. Through the ACL data collection initiative, some of the participants are making test sets available for tasks such as document retrieval and text categorization as well as knowledge and lexical resources. The ability to share resources and compare results on common tasks is a sign of the problem-oriented nature of TBIS research as well as the maturation of text processing and retrieval technologies.

—*Paul S. Jacobs*
General Electric Corporation

1991 AAAI Events Deadlines

October 15
Workshop proposals for AAAI-91 due. Contact Ed Lafferty (ell@mbunix.mitre.org)

November 1
Tutorial proposals for AAAI-91 due. Contact Elaine Rich (ai.rich@mcc.com)

November 16
Submissions for Spring Symposium Series due. Contact organizers listed below:
Argumentation and Belief. Sergio Alvarado (alvarado@iris.ucdavis.edu)
Logical Formalizations of Commonsense Reasoning. Vladimir Lifschitz (val@cs.stanford.edu)
Design of Composite Systems. Lewis Johnson (johnson@isi.edu)
Connectionist Natural Language Processing. Charles Dolan (cpd@aic.hrl.hac.com)
Constraint-Based Reasoning. Rina Dechter (dechter@cs.ucla.edu) or Sanjay Mittal (mittal.pa@xerox.com)
Implemented Knowledge Representation and Reasoning Systems. Charles Rich (rich@ai.mit.edu)
Integrated Intelligent Architectures. John Laird (laird@caen.engin.umich.edu)
Machine Learning of Natural Language & Ontology. Larry Reeker (reeker@cs.ida.org)

January 18 — Innovative Applications of Artificial Intelligence (IAAI) paper submissions due. Send to AAAI

January 30 — National Conference on Artificial Intelligence (NCAI) paper submissions due. Send to AAAI

February 2 — AAAI-91 brochure mailed. Allow three weeks for delivery

March 22 — IAAI Brochure and Preliminary Program mailed. Allow three weeks for delivery

March 26-28 — Spring Symposium Series, Stanford University

April 5 — AAAI-91 preliminary conference program mailed

June 15 — AAAI-91 and IAAI 91 housing deadlines

July 4-19 — AAAI-91, Anaheim, California

July 15-17 — IAAI, Anaheim, California