AI Magazine Volume 14 Number 3 (1993) (© AAAI)

AAAI 1993 Spring Symposium Series Reports

The American Association for Artificial Intelligence (AAAI) held its 1993 Spring Symposium Series on March 23–25 at Stanford University. This article contains summaries of the eight symposia that were conducted: AI and Creativity, AI and NP-Hard Problems, Building Lexicons for Machine Translation, Case-Based Reasoning and Information Retrieval, Foundations of Automatic Planning, Innovative Applications of Massive Parallelism, Reasoning about Mental States, and Training Issues in Incremental Learning.

Technical reports of the symposia AI and Creativity, Building Lexicons for Machine Translation, Case-Based Reasoning and Information Retrieval, Foundations of Automatic Planning, Innovative Applications of Massive Parallelism, Reasoning about Mental States, and Training Issues in Incremental Learning are available from AAAI. Instructions and an order form for purchasing electronic and hardcopy versions can be found elsewhere in this issue.

AI and Creativity

The symposium AI and Creativity attracted participants from widely differing backgrounds, including philosophy, science, education, engineering, and even computer science. The major themes of the meeting were the nature of creativity, computational models of creativity, and computational support for creativity. Despite their different backgrounds and research interests, participants interacted and communicated to an unusual degree, which was enormously exciting and satisfying.

The richness of the subject produced enthusiastic discussions but no simple solutions. On many issues, what was obvious to one participant was not at all clear to another. Even so, participants were enriched by the exchange of ideas and, at times, even changed their points of view as a result.

After much debate, the participants agreed on the following: (1) creativity is a multifaceted phenomenon, admitting different points of view from various disciplines; (2) judgments of creativity depend on the experience and knowledge of the problem solver; (3) creativity is a matter of degree rather than an all-or-nothing affair; (4) personal creativity, at the level of the individual, is a promising and manageable topic of study; and (5) deep knowledge is critical for creativity.

One topic that generated a lot of discussion was the role of representational redescription in creativity. Some participants argued that the redescription of a problem in a different representation can assist problem solving. Indeed, the judicious selection of a good representation can be highly creative: It can simplify the problem or even engender a solution that would otherwise have been unattainable.

The participants identified the following questions that still bring about much contention and could represent areas where future research could be directed:

First, should the goal of research in the area be to understand the creative process in humans, enhance it in humans, or produce creative machines?

Second, do we need to understand cognition before we can understand creativity, or does studying creativity help us to understand cognition?

Third, is the study of microdomains or toy worlds helpful for elucidating human creativity, or are such domains too simple and sparse for this purpose?

Fourth, is it useful to distinguish between creativity at the societal level (dealing with innovative results not previously known to the society) and creativity at the personal level (dealing with results that are new for the individual)?

Fifth, what is the role of experience in creativity?

Sixth, is the identification of the proper problem—as opposed to its resolution—the key to creativity?

The participants generally considered creative thought to be a subset of ordinary problem solving. Creativity is a special category, to be sure, often exhibiting distinct characteristics from those of routine decision making. Nevertheless, it appears to be similar in a qualitative sense to other modes of cognition.

Steven Kim
Lightwell, Inc.

Terry Dartnall
Griffith University (Australia)

Fay Sudweeks
University of Sydney (Australia)

AI and NP-Hard Problems

Numerous problems that arise in knowledge representation, planning, learning, and other areas of AI have been shown to be intractable. Although intractability can often be ignored on small, "toy" problems, it becomes a serious issue when one attempts to scale up AI techniques. This symposium was devoted to the fast-growing community of researchers working to understand, sidestep, or otherwise overcome intractability in AI problems.

There were two main threads running through the symposium: strategies for trading among the conflicting demands of generality, correctness, and computational complexity and methods for evaluating and understanding how well different solution methods address these conflicting demands. These two threads were evident in the opening talk by David Johnson, who discussed recent work on approximation methods for NP-complete problems and presented a list of "blind alleys" in research on NP completeness. He suggested that one such blind alley is the *classification disease,* where researchers produce increasingly large tabulations of how combinations of small variations on a problem affect computational complexity. Johnson also discussed the *small instance* and the *unrealistic distribution* blind alleys, which are characterized by experiments on problems that are too small or too unusual to give meaningful results.

Many of the papers emphasized the first thread, describing methods for dealing with intractability by sacrificing correctness or optimality. The topics ranged from approximation methods based on branch-and-bound to local search techniques (which can often find solutions even when the search space is too large to be searched exhaustively) to the handling of intractability when applying the CLASSIC knowledge representation system to real problems.

In some cases, approximate solutions are unacceptable, and the only known option is to design algorithms that are as efficient as possible (understanding that it might still require exponential time in some cases). Because such solution methods usually involve search processes, most of the symposium papers in this area were concerned with the development of efficient search methods. One of the papers that spurred the most debate was Matt Ginsberg's report on *dynamic backtracking,* an approach that stakes out a middle ground between traditional depth-first search and search techniques that maintain labelings in the style of assumption-based truth maintenance systems.

Many papers, either implicitly or explicitly, addressed the second thread: evaluation criteria for solution methods. In addition to well-established computational-complexity analysis methods, a body of recent work has investigated experimental approaches to evaluation. For example, a number of symposium papers explored the difficulty of propositional satisfiability problems generated according to a distribution suggested in the paper presented by David Mitchell, Bart Selman, and Hector Levesque at the 1992 National Conference on Artificial Intelligence (Hard and Easy Distributions of SAT Problems, pp. 459–465). These results add to a growing body of work suggesting that the hardest problems occur at transition regions in a problem space, such as the point where propositional satisfiability problems move from being mostly satisfiable to mostly unsatisfiable. This work sparked lively debate on the utility of experimental evaluations on randomly generated problems and, more generally, the difficulty of obtaining real problems for evaluating solution methods.

James Crawford
AT&T Bell Laboratories

Haym Hirsh
Rutgers University

Case-Based Reasoning and Information Retrieval: Exploring the Opportunities for Technology Sharing

This symposium was an attempt to begin to bridge the gap between practitioners in case-based reasoning (CBR) and information retrieval (IR). The CBR and IR communities have had little contact over the past decade, in spite of their mutual interest in the indexing, representation, and retrieval of online information. The symposium was structured as a set of invited talks to present the viewpoints of leading practitioners on their respective fields; paper presentations of work that either explicitly or implicitly brought concepts from one field to bear on problems typically associated with the other field were intermixed. This format worked well to draw out the differences and similarities between IR and CBR tasks as well as expose the mind sets that form the cultural background for the respective research efforts.

Bruce Croft began the symposium by comparing the goals of IR and CBR, declaring that the pragmatics of indexing large volumes of textual information has steered IR toward engineering solutions rather than cognitive modeling. IR has tended to be an empirical science, in which goodness is evaluated in terms of the precision and recall of real systems. However, the field lacks a good model of text representation, which could identify which features of text are most useful for retrieval in various situations. Janet Kolodner responded that it is precisely this lack of a good model of text representation that CBR research seeks to overcome. Because important features might not be evident in the surface representation of a case (or text), some prior interpretation is essential. Indexes need to be concrete enough to be recognizable in new situations and abstract enough to be broadly applicable. Indexes should be consistent with the world view of users and should fit the task that users are engaged in while asking for information.

Attendees debated the question of the potential incompatibility of IR's pragmatic approaches with CBR's goal of cognitively sound indexes. The conclusion was that the approaches were, for the most part, complementary, perhaps even forming a continuum. David Waltz argued that IR and statistical techniques allow one to make use of the raw data in unanticipated ways, finding patterns of data in large databases and making subsets of databases that are of a size amenable to analysis by case-based techniques. Edwina Rissland offered legal reasoning as a domain in which the retrieval of on-point cases depends on complex, goal-driven indexing and search strategies. The majority of applications, however, might well fall somewhere in the middle and benefit from a judicious mix of controlled and uncontrolled vocabulary indexing.

Workshop participants shared research efforts encompassing CBR for visual information and temporal sequences, extraction of structured information from text, knowledge base organization using conversational hypertext, and case-classification methods. These papers highlighted the multiplicity of approaches for helping users find information relevant to some task, and participants explored the viability of constructing hybrid systems.

As expected, the symposium ended with more questions than answers. It was generally agreed that bridging the current gap between huge text databases and relatively sparse case bases would also require help from the natural language and machine-learning communities. As the IR field moves toward multimedia support, its reliance on text-based methods will likely have to be modified to encompass structured representations of the sort that CBR researchers are pursuing, opening a new avenue for collaboration. Likewise, the statistical methods and relevance feedback approaches developed in IR might help CBR as it scales up and experiments with new modes of user-system interaction.

Peter Anick
Digital Equipment Corporation

Evangelos Simoudis
Lockheed Corporation

Foundations of Automatic Planning: The Classical Approach and Beyond

Over the past 10 years, AI planning has experienced dramatic diversification and growth. Early work led to a natural evolution of classical planners, all based on the modal-truth criterion. During the late 1980s, however, planning researchers diverged from this traditional trend, focusing on methods for coping with uncertainty and dynamic domains. These forays into reactive planning took on many forms, including procedural and case-based reasoning, types of program compilation, and extensive state-based reasoning compiled into reactively applicable control rules. Interestingly, the 1990s have seen a renewed focus on classical planning. Most of this research emphasizes theoretical analysis and empirical studies of the basic techniques of traditional planning. This new trend led to much discussion at the symposium.

The symposium's goal was to take stock of AI planning, to foster an introspective analysis of the field in general and classical planning's foundational role in particular. Organizers hoped to encourage discussion that constructively analyzed the inherent strengths and weaknesses of the classical approach. Another goal was to relate new formalisms and approaches to classical planning. Finally, they hoped to establish more solid agreement on planning vocabulary, clarifying the status of current results and outstanding problems.

The symposium got off to an exciting start with an invited talk by David Wilkins. He set the stage for a controversy, which lasted throughout the meeting, about the kind of theoretical-empirical studies that have proliferated in the field. Wilkins seriously questioned the basic utility of recent planning work. He argued that builders of planning systems designed for application to realistic domains know that the "vanilla" modal-truth–criterion formulation is inadequate. For example, it ignores many important features, such as task-reduction schemas and critics, that are intrinsic even to traditional planners such as NONLIN and SIPE. Despite the practicality of such features, they have not been considered by more recent analytic and empirical studies. It was widely felt that this gap must be bridged; planning researchers must try to ensure that their abstractions and assumptions are reasonable in terms of the requirements of realistic problems. Otherwise, the results and insights gained by such studies can only be of limited use to practical planning systems. Several people joined Wilkins in lamenting the ironic effect that this trend has had on the ability to publish. Just as the focus on reactivity and other real-world issues made it difficult to publish theoretical-empirical work in the late 1980s, the renewed focus on classical planning has created a dearth of forums for work on implemented systems that incorporate new techniques or focus on novel planning applications.

Two of the four symposium panels echoed these tensions: one on the building of real planners and the other on the role of formalization. A third panel focused on alternative planning methodologies and underscored the need for integrating and understanding all planning methods as a collective whole. Drew McDermott echoed this theme of planning diversity, describing his system that strives to integrate reasoning about behavior with a highly expressive procedural framework for robot programming. The fourth panel dealt with the role of empirical evaluation—how researchers should perform their studies and how they should use well-principled evaluation to guide their development of new ideas. Organizers feel this meeting can play a pivotal role in shaping the future of the planning field, a task that will be taken up again in 1994 at the biannual planning conference.

Amy Lansky
Mark Drummond
NASA Ames Research Center

Subbarao Kambhampati
Arizona State University

Ed Pednault
AT&T Bell Laboratories

Qiang Yang
University of Waterloo

Building Lexicons for Machine Translation

The lexicon plays a central role in machine translation (MT), but it is only recently that MT researchers have begun to focus specifically on the problem of building large dictionaries. This task is especially challenging for MT because of cross-linguistic divergences and mismatches that must be handled during the translation process. The symposium Building Lexicons for Machine Translation brought together researchers from the fields of MT and the lexicon to discuss relevant issues.

One of the most entertaining moments of the meeting came when Sergei Nirenburg gave his 10-minute overview of the symposium during the plenary session of the symposium series. He stated that the burning issues raised by the MT practitioners were the following:

Content: How little can we get away with?

Form: How dare you say that your formalism is better than mine?

Provenance: Wouldn't it be nice if somebody else did it for us?

Clients: Must we really be bothered by all these messy phenomena?

No doubt one of the reasons Nirenburg raised the content issue was that he noted several presentations where the lexical representations were encoded minimally. These researchers did not invoke a cry for minimalism on theoretical grounds but, rather, implicitly on practical grounds.

The topic of preferred formalisms arose following a number of presentations on formal mechanisms, including one by Ann Copestake and Antonio Sanfilippo that proposed a mechanism for linking typed feature structures in lexical entries. When Christian Rohrer questioned the justification for the typed formalism over nontyped feature systems, Copestake pointed out that typed features facilitate the identification of errors.

Nirenburg's third issue is a curious variation on the classic divide-and-conquer strategy; he suggests that some MT researchers are splitting up the work to avoid conquering the tough problems. Perhaps it was this observation that led to Lori Levin and Nirenburg's presentation, in which they described a continuum of translation cases that extends from those that can be handled by predictable principles to those that must be handled in an idiosyncratic way.

Probing further into these unruly cases, two linguists not widely known in the MT community gave invited talks about their work in lexical semantics. Beth Levin presented her working hypothesis that syntactic behaviors are generally determined semantically. Alan Cruse presented his theory of polysemy and discussed the inadequacies in the qualia theory of James Pustejovsky.

The panels organized by Scott Bennett and Nirenburg, as well as the invited talk by Makoto Nagao, provided a forum for discussion about Nirenburg's fourth issue, namely, those messy user concerns that are often ignored by MT researchers. Bennett pointed out two examples: the need to develop MT systems that are robust enough to cope with newly coined terms and phrases and the need to let users do lexicon customizing themselves.

The symposium ended with an announcement by Susann Luperfoy about the COMLEX project. Part of the linguistic data consortium at the University of Pennsylvania, the COMLEX project seeks input from the MT community on defining what belongs in the dictionary. The purpose of this project is to provide a common, sharable lexicon for speech and language technology.

Bonnie J. Dorr
University of Maryland

Innovative Applications of Massive Parallelism

Massively parallel AI as a paradigm is driven by a number of underlying assumptions. Pragmatic computational considerations, neurocognitive considerations, and computational infrastructure considerations can be identified as common ground for most of the researchers in this area.

Pragmatic computational considerations center on the fact that even a small quantitative change in the speed of performing a task often results in dramatic qualitative differences. This view is supported by many areas of real-life endeavor but also by elementary physics, such as in the case of the escape velocity from a planet. This basic organizing principle has been obfuscated in knowledge representation by the mainstream focus on computational complexity issues. This focus has led many researchers to neglect implementational issues and stress results of what cannot be done under worst-case conditions.

Neurocognitive considerations focus on well-known, basic observations about the human cognitive-processing system, such as Jerry Feldman's 100-step rule. This rule states that many of the interesting human cognitive abilities can be performed in a time frame of about 500 milliseconds, which forces the use of many parallel components, each of which can perform at most 100 operations for one task.

Computational infrastructure considerations are based on the historically unique development of computational hardware. Computational hardware has improved in power, speed, and availability and simultaneously become many orders of magnitude smaller and cheaper. Because traditional supercomputers have reached limits dictated by elementary physics, the current revolution in computational performance advancement is carried by the massively parallel generation of computers.

With the availability of large amounts of cheap processing power, a reorientation of the important parameters of computer use becomes necessary. This reorientation seems more acceptable to researchers in massively parallel AI than to mainstream computer scientists or classical AI researchers. Looking at the history of AI, we find precedents for progress in developing better hardware.

Those attending the symposium spanned an amazing range of different work areas and paradigms. Although AI and neural network researchers seem to have stopped communication in most other academic settings, a good mix of symbolic and connectionist approaches was represented here. Papers came from such diverse areas as search, natural language, applications, actors, knowledge representation, neural networks, constraint satisfaction, logical reasoning, objects, planning, genetic algorithms, rule-based systems, machine learning, and hardware design.

The large number of papers in the search category indicates that parallel AI might, to some degree, go through a similar development as AI itself did, having many of its roots in search algorithms. The papers in areas such as knowledge representation and machine learning, however, raise the hope that the developmental cycle of the field will be much shorter.

Perhaps the most important parts of this symposium were two panels, where questions about the future research agenda were repeatedly raised. Three statements about the field were established: First, we in the field are bullish about massive parallelism and AI. Second, massive parallelism will not supplant any other AI fields but hopefully will become as useful and natural to all of them as the computer has become to cognitive science. Third, we have to make a massive effort to avoid the unreasonable expectations and hype that have resulted in so much disappointment with AI.

James Geller
New Jersey Institute of Technology

Reasoning about Mental States: Formal Theories and Applications

Researchers in AI often reason informally in a way that involves attributing a variety of mental states to machines and, often, design systems that must be able to reason about both their own mental states and those of others. This kind of reasoning is common in a number of areas, including cooperative interfaces for databases; database security; planning; and, especially, multiagent planning.

Until recently, most formal work within AI on mental states has concentrated only on a related pair of notions—knowledge and belief. In the past few years, however, an increasing reliance on a wide variety of mental notions in the design and understanding of actual systems has led to a broadening of this formal work. As a result, formal theories now exist within AI about a number of mental notions and their close relatives, including ability, action, choice, commitment, desire, intention, goals, obligation, and perception.

Our aim was to bring together researchers working on formalisms for reasoning about these mental notions as well as researchers involved in the design of systems that rely on or incorporate these notions. The symposium began with an invited talk by John McCarthy, reviewing some ideas from his early paper "Ascribing Mental Qualities to Machines," and included an invited presentation by Michael Bratman on recent philosophical work on the notion of intention. Apart from these talks, the symposium was organized around contributed papers, with a few general discussion periods.

The range of contributed papers was rather broad. Quite a few papers were on the proper logical understanding of intentions, in keeping with the recent interest in the foundations of belief-desire-intention–style architectures. The symposium also included talks on the problem of incorporating several distinct attitudes within a single logical framework; the problem of updating various attitudes, not just knowledge and belief; the notion of action and ability, both from a model theoretic and an algorithmic point of view; a logical treatment of perception; the development of a general logical theory of practical reasoning; and the notion of obligation and other topics from deontologic. In addition, attendees heard descriptions of several systems that incorporate a notion of mental states and at least one criticism of the relevance of the formal work devoted to these notions because of the kind of understanding of mental states that is needed in designing actual systems.

The topic of this symposium is a relatively new area of interest within AI. Many of the researchers working within this area were unacquainted with the work of others and were surprised to learn of work along similar lines. We feel that the symposium was successful in helping to focus research on the development of precise theories for reasoning about a variety of mental states and also that it will serve to stimulate interaction between those whose research relies on these theories and those who are concerned primarily with the logic of the matter.

John Horty
University of Maryland

Yoav Shoham
Stanford University

Training Issues in Incremental Learning

In most realistic learning settings, data are presented in streams, there are time and complexity constraints on agents, target concepts might drift upon time, and the learning task and success criteria can vary. Clearly, these requirements call for some sort of incremental variations in learning.

What really is incremental learning? What makes it different from nonincremental or direct learning? Although many researchers have studied incremental learning systems, until recently, little effort has been spent identifying and studying issues specific to incremental learning.

One topic of discussion at the symposium was just what is a good definition of incremental learning. The behavioral definitions used until now—for example, a learning algorithm is incremental if it inputs one training instance at a time—are inadequate for distinguishing between incremental and nonincremental learning. One also needs to forbid the possibility of reprocessing all past observations. The consensus at the symposium was that learning is *incremental* when, in addition to the earlier constraints, there are memory-size limitations that prevent the learner from encoding all the data. It appears that the theoretical community, to which approximately half the participants belonged, has already realized significant advances in the area of space-bounded learning, yielding various bounds for different assumptions on memory limitations. Thus, there were lively and fruitful exchanges between the theorists and the empiricists, members of both communities eager to learn more about the questions and results of mutual interest.

The second big topic of discussion was ordering effects. Because of their constraints, incremental learning systems are subject to ordering effects; that is, they can reach different states for different training sequences over the same collection of data. This property is, in fact, one of the main reasons that the study of incremental learning is worthwhile. If it were not for this effect, incremental learning would differ from nonincremental learning only in terms of computational complexity. Ordering effects have been mentioned widely in the literature but almost never addressed specifically. At the symposium, the cause, avoidance, desirability, and utility of ordering effects were all debated. The symposium initiated a wide interest in this topic that should grow in the future.

Finally, given that revision of hypotheses is central to incremental learning, it was found that strong links exist with all approaches that aim to evaluate different competing hypotheses and seek ways to propose new ones. In fact, one invited lecturer argued convincingly that nonmonotonic reasoning concepts and methods could possibly offer a natural framework in which to study incremental learning.

To sum up, the symposium helped to determine and clarify issues that are important on the research agenda. Most importantly, the openness of everyone and the congenial atmosphere that resulted helped foster new collaborative efforts between people from different backgrounds.

Antoine Cornuejols
University of Paris at Orsay

Douglas Fisher
Vanderbilt University

Sally Goldman
Washington University

Lorenza Saitta
Universita di Torino

Jeffrey Schlimmer
Washington State University