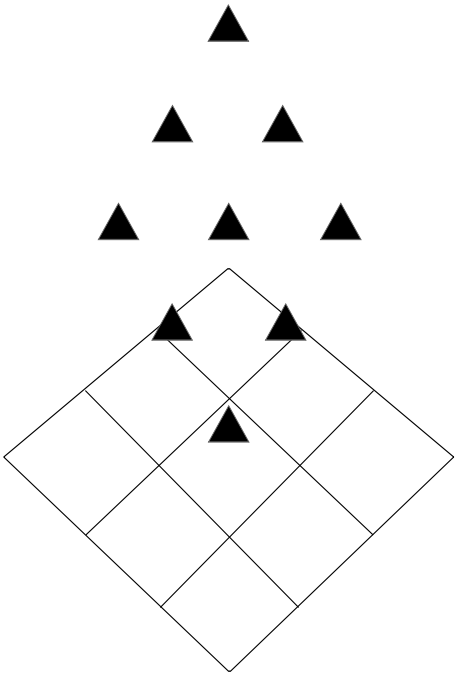

AAAI Spring Symposium Series

Call for Participation

March 23, 24, 25, 1993

Stanford University
Stanford, California

The American Association for Artificial Intelligence presents the 1993 Spring Symposium Series, to be held Tuesday through Thursday, March 23–25 at Stanford University, Stanford California.

The topics of the eight symposia in the 1993 Spring Symposium Series are:

- AI and Creativity
- AI and NP-Hard Problems
- Building Lexicons for Machine Translation
- Case-Based Reasoning and Information Retrieval—Exploring the Opportunities for Technology Sharing
- Foundations of Automatic Planning: the Classical Approach and Beyond
- Innovative Applications of Massive Parallelism
- Reasoning About Mental States: Formal Theories and Applications
- Training Issues in Incremental Learning

Most symposia will be limited to between forty and sixty participants. Each participant will be expected to attend a single symposium. Working notes will be prepared and distributed to participants in each symposium.

A general plenary session, in which the highlights of each symposium will be presented, will be held on Wednesday, March 24, and an informal reception will be held on Tuesday evening, March 23.

In addition to invited participants, a limited number of other interested parties will be able to register in each symposium on a first-come, first-served basis. Registration will be available by December 1992. To obtain registration information write to the AAAI at 445 Burgess Drive, Menlo Park, CA 94025 (sss@aaai.org).

Submission Requirements

Submission requirements vary with each symposium, and are listed in the descriptions of the symposia. Please send your submissions directly to the address given in the description. Do *not* send submissions to AAAI. All submissions must arrive by October 16, 1992. Acceptances will be mailed by November 16, 1992. Material for inclusion in the working notes of the symposia will be required by February 1, 1993.

AI & Creativity

The symposium will provide a forum for exploring current research in, or relevant to, artificial intelligence and cognitive science which pertains to creativity. It will draw together researchers from diverse disciplines as well as practitioners and managers engaged in projects that require creativity.

Relevant disciplines include AI, cognitive science, computer science, psychology, philosophy, logic, design, engineering, and others.

Topics include, but are not limited to philosophical issues in the computational study of creativity; psychological theories and models of creativity; conceptual and formal models of creativity; representational change and redescription; analogy and abduction; theory generation and discovery; hybrid architectures (e.g., classical and connectionist hybrids; hybrids using genetic algorithms to evolve neural nets; interplay of deterministic and probabilistic strategies); use of AI and other computational techniques to enhance creativity; and application areas (e.g. science, engineering, architecture and design, business, government)

The symposium will maintain a balance between theoretical issues and descriptions of implemented systems to promote synergism between theory and practice.

Authors should submit extended abstracts of 2000 words. These should be sent by email to Fay Sudweeks (fay@archsci.arch.su.oz.au), in conjunction with a hard copy by October 16, 1992. For further information contact:

Fay Sudweeks, Organizer (fay@archsci.arch.su.oz.au)
AAAI Spring Symposium on Creativity
Department of Architectural and Design Science
University of Sydney, NSW 2006 Australia

Program Committee: Terry Dartnall (cochair), Steven Kim (cochair) Robert Levinson, Fay Sudweeks, Devika Subramanian

AI and NP-Hard Problems

Numerous problems that arise naturally in knowledge representation, learning, planning, and other areas of AI can be shown to be NP-hard. Many of these problems involve computationally similar issues cast in different domain-specific guises. This symposium is devoted to the fast-growing community of researchers studying the relationship of AI problems and algorithms to more traditional problems and algorithms from the theoretical computer-science community.

For example, stripping away domain-specific details and applying AI methods directly to intractable problems in their "pure'' form (e.g., on problems formulated in terms of mathematical primitives such as graphs and sets) makes it possible to generalize results and minimizes the potential for repeatedly attacking the same problem in different domain-specific guises. Moreover, such "abstracted'' versions of AI problems provide a good testbed for comparing the behavior of competing algorithms.

Relevant work should be grounded in actual algorithms (ideally with implementations), or should make clear contributions to such efforts. The symposium will be organized in a format that combines research presentations with open-ended discussions. Each paper session will have a number of short paper presentations, followed by a critical review designed to generate discussion. Those who wish to make presentations should submit a draft paper of length at most ten pages or alternatively, an extended abstract of at most five pages. All such submissions should include a list of keywords to aid in the reviewing process. Those who wish to attend without presenting a paper should submit a description of their research interests and a list of relevant publications.

Four copies of all submissions should be sent to arrive by October 16, 1992 to:

Haym Hirsh (hirsh@cs.rutgers.edu)
Department of Computer Science
Hill Center for the Mathematical Sciences
Busch Campus, Rutgers University
New Brunswick, NJ 08903
Fax: 908-932-5691 / Phone: 908-932-4176/2001

Program Committee: James Crawford, Rina Dechter, Tom Ellman, Haym Hirsh (chair), David McAllester, Steve Minton, Bart Selman

Building Lexicons for Machine Translation

The lexicon plays a central role in any machine translation (MT) system, regardless of the theoretical foundations upon which the system is based. However, it is only recently that MT researchers have begun to focus more specifically on issues that concern the lexicon, e.g., the automatic construction of multi-lingual semantic representations. Large dictionaries are important in any natural language application, but the problem is especially difficult for MT because of cross-linguistic divergences and mismatches that arise from the perspective of the lexicon. Furthermore, scaling up dictionaries is an essential requirement for MT that can no longer be dismissed.

This symposium provides a forum for researchers from the fields of MT and the lexicon focus on the intersection of the two fields, rather than their broader concerns. A number of fundamental questions will be addressed:

What lexical levels are required by a machine translation system? Syntactic? Lexical semantic? Ontological? What do the representations at each of these levels look like, and how would they be constructed?

What are the interdependencies between these levels? Can we take advantage of interacting linguistic constraints from each level for the construction of lexical representations? Should the levels be kept as separate layers and related explicitly or should they be combined into one layer and be related implicitly? Should all levels be represented in the same or in different, dedicated formalisms? What are the implications of these choices for MT system architecture, processing of the relevant knowledge, interaction between components of MT systems, applicability of the resulting knowledge sources in different types of MT mappings?

Can automatic procedures be used for the construction of lexical representations? What existing resources should we be using and what aids do we have to transform these resources into appropriate representations for MT? To what extent is it possible to acquire elements of contrastive knowledge (mapping information) using existing techniques (e.g., work on bilingual corpora, example based approaches, etc.)?

To what extent is it possible to share lexicons? If the representations and the actual knowledge are tailored to a specific system (e.g., style of grammar or choice of domain knowledge base) then how can sharing be achieved? How much representations and knowledge are tied to specific approaches to MT system construction, and, to the extent that they are, how much can people come to some agreement on some of those other issues so that they can share lexicons?

Are bilingual dictionaries useful for the construction of computational lexicons for MT? What is the role of example sentences and phrases in bilingual dictionaries? Can we extract information from pairwise examples in order to achieve example-based translation? Can we use bilingual dictionaries for the extraction of grammatical information?

What are the different types of MT mappings (transfer, interlingual, statistically based, memory-based, etc.) and how do these mappings affect the representation that is used in the lexicon?

What types of MT divergences and mismatches must be accommodated in the lexicon (i.e., cases where the target-language sentence has a different structure, or conveys different information, from that of the source language)? Are these problems that any translation system must deal with regardless of the MT mapping that is used? If so, can we construct lexicons that accommodate these divergences regardless of the translation mapping that is used? Can we incorporate information about the respective portions of lexical/non-lexical knowledge needed to decide on suitable candidates for target constructions and on lexical clues for strategies for such decisions?

Can we, or have we, achieved language-independence in the representations that are used in the lexicon? Can we support an interlingual approach to machine translation based on current technology and resources?

All interested participants should submit five copies of a one- to five-page abstract (not including the bibliography) by October 16, 1992 to:

Bonnie Dorr (bonnie@umiacs.umd.edu)
Department of Computer Science / UMIACS
University of Maryland
A.V. Williams Building
College Park, MD 20742

Fax or electronic submission will not be accepted. Each submission should include the

names and complete addresses of all authors. Correspondence will be sent to the authors by e-mail, unless otherwise indicated. Also, authors should indicate under the title which of the questions and/or topic listed above best describes their paper (if none is appropriate, please give a set of keywords that best describe the topic of the paper). Submissions will be judged on clarity, significance, and originality. An important criterion for acceptance is that the abstract clearly contributes to the theme of building lexicons for machine translation. Abstracts focusing on one of these two areas (i.e., MT or the lexicon) will be given a lower priority than those that address issues that lie at their intersection.

Program Committee: Michael Brent, Bonnie Dorr (chair), Sergei Nirenburg, Elaine Rich, Patrick Saint-Dizier

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

The fields of case-based reasoning (CBR) and information retrieval (IR) have a shared interest in searching databases to locate information that is relevant to a problem being solved, and using the information to shed light on or solve the problem. As such, both fields have shared interests on: indexing of information, the formulation of query expressions suited to retrieving relevant information, heuristic matching, the measurement of similarity, and the use of domain knowledge to improve search. Whereas case-based reasoning research has historically worked with small collections of well structured cases requiring a fair degree of hand-tailoring, researchers in information retrieval have concentrated on indexing and querying over very large collections of primarily textual data, with the aim of minimizing the need for hand-tailoring. IR systems typically present to people the articles retrieved in response to a query, who then read, use, and discard them as appropriate. On the contrary, case retrievers are often set in a larger automatic problem-solving system that frequently includes a learning component allowing the dynamic reorganization of the case base. Perhaps due to these differences, there has been relatively little interaction to date among practitioners in the respective fields.

Case-based reasoning resarchers are now embarking on an ambitious second phase for their work whose goal is the implementation of systems that use large databases containing a variety of information types. For example, Cognitive Systems Inc has applied CBR to classify the text of telex messages. A case base of several thousand thematically cross-indexed cases (in video form), for use in story-based teaching, is currently under development at the Institute for the Learning Sciences.

Knowledge-based approaches have been introduced into information retrieval systems as well, in the form of thesauri, semantic nets, concept frames, etc. There is a recent trend in the IR community toward extracting and utilizing structured information to complement full-text retrieval methods, and to extend textual retrieval systems to encompass multi-media.

Not only does the intersection of the two fields appear to be growing rapidly, but many tasks (such as text categorization) and numerous application domains (such as legal, medical, help-desk) may require both IR and CBR, raising the question of how to best integrate the methodologies into a single system with a uniform interface. Finally, an increasing number of databases contain mixed text and non-text data, and can draw on and benefit from both methodologies.

The purpose of this symposium is to bring researchers from both communities together to discuss issues of common interest, share the results and experiences of their respective research, and seek areas of potential future technology transfer or convergence.

Specific topics of interest include, but are not limited to, such questions as:

How might IR integrate and take advantage of more structured information, as used in CBR? Can weak and strong retrieval methods be effectively combined?

How does CBR scale up to large collections of semi-structured information? How can CBR minimize the need for hand-tailoring of the data in a case base?

What similarity assessment methods and metrics have been developed in CBR and IR? What demands do they make on knowledge representation schemes?

How well do CBR/IR techniques apply to multi-media information bases?

How can the construction of viable queries for retrieving desired information be facilitated through CBR methods?

What kinds of knowledge representations are needed to support reasoning (over cases) as opposed to retrieval? What is the role of reasoning in retrieval?

Can/should the functions of textual information retrieval and case-based reasoning be integrated into a single application? How can the effectiveness of such a hybrid CBR/IR system be evaluated?

The symposium will consist of individual presentations and panel discussions with ample opportunity for group discussion. We will strive, as the first priority, for equal participation from the two communities, and secondarily for a balance between academia and industry.

Those who wish to make presentations should submit a draft paper, of length at most ten pages. All prospective submitters are encouraged to contact Evangelos Simoudis or Peter Anick to discuss how what they wish to present is coordinated with the objectives of the symposium. Those who wish to attend without presenting a paper should send a description of their research interests and a list of related publications. Four copies of all submissions should be sent to arrive by October 16 to:

Evangelos Simoudis
Lockheed AI Center
O/96-20 B/254F
3251 Hanover Street
Palo Alto, CA 94304
simoudis@titan.rdd.lmsc.lockheed.co

Program Committee: Peter Anick (cochair), Bruce Croft, William Mark, Chris Riesbeck, Evangelos Simoudis (cochair)

Foundations of Automatic Planning: The Classical Approach and Beyond

The focus of this symposium will be to take stock of where AI planning has been and where it is going—an "introspective analysis" of the field in general and classical planning's foundational role in particular. The symposium will foster discussion on various fronts. First of all, we hope to constructively analyze inherent strengths and weaknesses of the classical approach. While previous classical planning systems may have failed to achieve certain types of behavior, the underlying precepts of classical planning may not themselves be flawed. Another goal is to relate new formalisms and approaches to previous work in classical planning. Finally, since terminology has often become confused in recent years, one of the goals of the symposium will be to establish a more solid agreement on a planning vocabulary. By doing so, we hope to clarify the status of current results and outstanding problems.

Papers are solicited in the following topic areas:

Techniques: This area includes algorithms, heuristics, planning representations, and planning strategies. We encourage participants to present their work in a way that fosters comparative understanding within the field, especially with respect to classical planning. We also encourage the submission of papers that attempt to summarize the state of the art in a particular area, or attempt to collect together disparate contributions under a single coherent view.

Terminology: This area will focus on helping to define a common language that researchers can use to more effectively communicate. We hope to provoke discussion on some of the more commonly abused terms and to form a better common understanding of the salient referents of planning terminology.

Research Methodology and Evaluation: This area will focus on methodologies for planning research and how work in planning can be evaluated and compared. One area of recent interest is the relative merits of utilizing toy domains, simulated domains, real-world domains, and domains situated within a real physical environment—i.e., a

focus on the qualities that domains should manifest and what domains tell us about planner performance. We also encourage the submission of papers that attempt to define metrics for the evaluation of AI planning systems and establish a vocabulary useful for the comparison of different planning methods. We also encourage the presentation and dissemination of new problems and domain data that can serve as common testbeds for researchers in the field.

Those wishing to participate should submit six copies of a short paper (approximately five pages) as well as a one-page statement of research interests and bibliography. Those only wishing to attend should submit six copies of a one-page research statement and bibliography. All submissions should include an email address, a telephone number, and a mailing address. Prospective participants are encouraged to contact members of the program committee (preferably by e-mail) with any questions or comments.

Submissions should be sent to arrive by October 16, 1992 to:

Amy Lansky—Attn: Spring Symposium 1993
NASA Ames Research Center
MS 269-2 Moffett Field, CA 94035-1000
lansky@ptolemy.arc.nasa.gov

Program Committee: Mark Drummond, Subbarao Kambhampati, Amy Lansky (chair), Ed Pednault, Qiang Yang

Innovative Applications of Massive Parallelism

Artificial intelligence, in many respects, has not met the expectations of the scientific, defense, and business communities. Frustration with the problems of the knowledge-based, symbolic approach that was firmly established in the AI community in the 1970s and early 1980s has lead some researchers to return to low-level connectionist and artificial neural approaches. It has become clear, though, that neither the symbolic nor the sub-symbolic approach will offer, any time soon, a complete solution to the problems of AI. The 1980s were also marked by the development of commercially viable new hardware designs that made use of massive parallelism. Such machines are characterized as having processor resources that are relatively abundant compared to the size of the target problem. Typically, the processors in massively parallel machines number in the thousands.

So far, experiences in this area indicate that both symbolic and connectionist approaches can benefit from massively parallel implementation. Symbolically oriented massively parallel AI implementations combine the elegance of established formalisms of fields such as knowledge representation and search with the brute force of the underlying hardware and the comfort of a fully developed programming environment. Neural network simulators can benefit from implementation on massively parallel hardware.

The purpose of this symposium is to bring together researchers in different areas of artificial intelligence who are working in a paradigm using massive parallelism. We are especially interested in massively parallel search, knowledge representation, information retrieval, actor systems, natural language processing, constraint satisfaction, genetic algorithms, expert systems, and neural network realizations. However, the symposium is open to any innovative approach to using massive parallelism for AI systems. The objective of the symposium is to exchange ideas and stimulate further research in how massively parallel computing can benefit AI.

To be considered for invitation to the symposium on Innovative Applications of Massive Parallelism, please submit either an extended abstract or a position statement on the subject of the symposium. Submissions should be ten pages or less (ten-point, double-spaced). Hardcopies should be sent to the chair of the symposium by October 16, 1992:

James Geller (geller@vienna.njit.edu)
New Jersey Institute of Technology
CIS Department, University Heights
Newark, NJ 07102

A cover letter with author(s), title, affiliation(s), contact address, phone, and email address should be included. Work in progress is acceptable. Abstracts of recently published papers or papers submitted concurrently to other conferences are acceptable but this fact should be mentioned in the cover letter.

Program Committee: Gul Agha, Matt Evett, James Geller (chair), Hiroaki Kitano, Curt Powley, Dave Waltz

Reasoning about Mental States: Formal Theories & Applications

Researchers in AI often design systems that must be able to reason both about 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, intelligent tutoring and especially multi-agent coordination. However, until recently, most formal work on mental states within AI 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 mentalistic notions in the design and understanding of actual systems has led to a broadening of this formal work. As a result, there now exist within AI formal theories of a number of mentalistic notions and their close relatives, including: ability, action, choice, commitment, desire, intention, goals, obligation, perception.

The aim of this symposium is to bring together researchers working on formalisms for reasoning about these mentalistic notions, and also researchers involved in the design of systems that rely upon or incorporate these notions. We do not wish to exclude innovative work on knowledge and belief, especially if it concerns connections with other mental states, but do explicitly want to emphasize a wider variety of mental states. We would expect papers on topics such as the following:

Descriptions of systems, architectures, or theories that rely explicitly upon reasoning with mentalistic notions (e.g., BDI-style planning architectures, multi-agent planners, models of discourse and speech acts);
Formalisms for reasoning about particular mentalistic notions within particular frameworks (e.g., a dynamic logic for reasoning about intention) and comparisons between frameworks (e.g., between representations of action in dynamic and temporal logics), or for reasoning about a number of these notions and their interactions (e.g., a logic for reasoning about intention, belief, obligation, and action);

Formalisms for multi-agent mentalistic reasoning and coordination (e.g., reasoning about the beliefs of one agent concerning the intentions and beliefs of another);

Realistic (non-idealized) models of these various mental states;

Issues involving the interaction between these various mentalistic notions and time (e.g., the temporal persistence of commitment);

Issues concerning the computational relevance of sophisticated philosophical theories of mental states, and perhaps also concerning the relevance of the recent, detailed models developed within AI for philosophical theorizing.

We hope that the symposium will help to focus research on the development of precise theories for reasoning about a variety of mental states, and also that it can serve as a forum for interaction between those working in other areas whose research relies upon these theories and those concerned primarily with the logic of the matter.

Those wishing to present their work for discussion at the symposium should submit five physical copies of an extended abstract, no more than six pages in length. (Authors whose work is accepted for presentation will later be asked for a full paper to be included in the symposium working notes.)

Those wishing to attend the symposium without presenting work should submit a one page statement of their (relevant) research interests.

All submissions and requests for attendance should be sent to the following address by October 16:

John Horty (horty@umiacs.umd.edu)
Inst. for Advanced Computer Studies
A.V. Williams Building, University of Maryland
College Park, MD 20742

Program Committee: Jon Doyle, John Horty (cochair), Hector Levesque, Martha Pollack, Yoav Shoham (cochair)

Training Issues in Incremental Learning

Unlike the batch learning model, in which a learner receives all the data before processing them, in the incremental (or on-line) learning model a learner must process the data as it arrives. Thus, an incremental learning algorithm must make predictions in an on-line manner, typically updating the hypothesis throughout the

learning session. While many researchers have studied incremental learning systems, until recently, little effort was spent identifying and studying issues specific to incremental learning. For example:

- How do the complexity of learning (cost of updating the hypothesis)and the improvement in the accuracy of the hypothesis change during the course of the learning session?
- How many observations are required to obtain a "stable" hypothesis?
- How does data ordering affect learning?
- How can concept drift be handled?

In addition to their practical implications, exploring these issues may uncover important connections with other fundamental issues in machine learning. For example, recent works suggest that deep relationships exist between data ordering, bias, and theoretical concepts such as teachability and information content. Accordingly, the focus of this symposium will not be on implementations, but on the general properties and common issues addressed by the many separate research efforts in the area of incremental learning. The symposium will provide an opportunity for researchers interested in this area to share their results and hopefully in the process gain a better understanding of the issues and techniques important to this growing domain.

Topics of interest include but are not restricted to: the relationship between data ordering and bias; the necessary and sufficient conditions (on the learner and data) to get order independence; the tractability of identifying a best order or teaching strategy; valiant-style theoretical treatments of the quality of results as a function of data or teaching strategy; and the study of techniques (that may vary over time) to determine which examples to ignore during the learning session.

The symposium will be organized in a format that combines research presentations with open-ended discussions. It will begin with brief introductory remarks posing a list of issues (prepared in advance by the organizing committee) to be addressed during the symposium. Papers will be organized into sessions according to the primary issues addressed. We expect each session to have roughly three paper presentations. For each session the committee will have selected a commentator familar with the area who will lead an open-discussion session about issues raised by the presentations and interesting research directions to build upon the results presented. The commentator will be given copies of the paper in advance, and may briefly summarize other work that is related to the papers presented in the session.

Those interested in presenting their work should submit *seven* copies of either a short paper (five-ten pages) describing completed work, or an extended abstract (two-five pages) describing preliminary work. Those interested only in participating should submit a one page summary of their research interest, together with a list of related publications.

Submissions should be sent to arrive by October 16, 1992 to :

Antoine Cornuejols (antoine@lri.lri.fr)
L.R.I., Universite de Paris-Sud
Bat. 490, 91405 ORSAY Cedex
FRANCE

Program Committee: Antoine Cornuejols (chair), Douglas Fisher, Sally Goldman, Lorenza Saitta, Jeffrey Schlimmer.