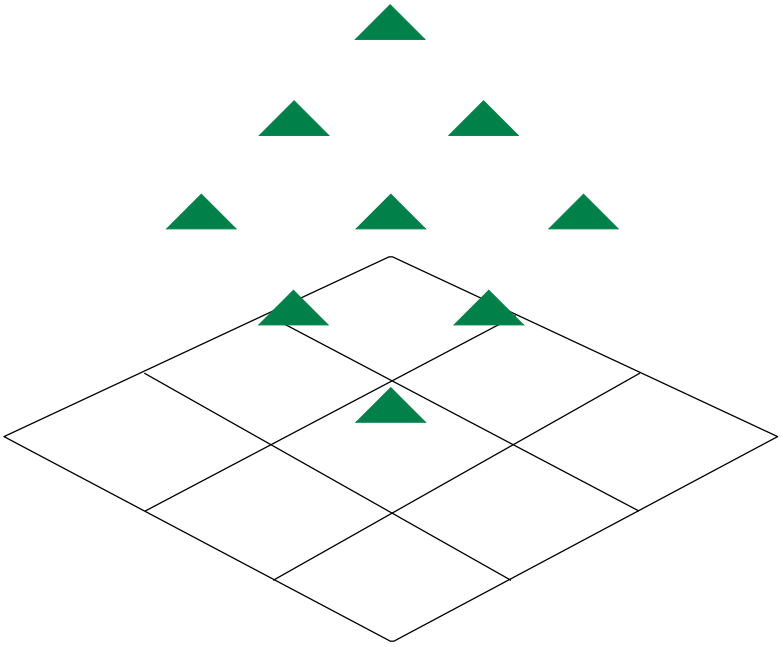

AAAI 1994

Fall Symposium Series

November 4–6, 1994
The Monteleone Hotel
New Orleans, Louisiana

Call for Participation

Sponsored by the
American Association for Artificial Intelligence
445 Burgess Drive, Menlo Park, CA 94025
(415) 328-3123
fss@aaai.org

The American Association for Artificial Intelligence presents the 1994 Fall Symposium Series, to be held Friday through Sunday, November 4–6 at the Monteleone Hotel, New Orleans, Louisiana. The topics of the five symposia in the 1994 Fall Symposium Series are:

- **Control of the Physical World by Intelligent Agents**
- **Knowledge Representation for Natural Language Processing in Implemented Systems**
- **Improving Instruction of Introductory AI**
- **Planning and Learning: On to Real Applications**
- **Relevance**

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 November 5, and an informal reception will be held on November 4.

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 mid-July 1994. To obtain registration information write to the AAAI at 445 Burgess Drive, Menlo Park, CA 94025 (fss@aaai.org).

Submission Dates

- Submissions for the symposia are due on April 15, 1994.
- Notification of acceptance will be given by May 17, 1994.
- Material to be included in the working notes of the symposium must be received by August 19, 1994.

See the appropriate section below for specific submission requirements for each symposium.

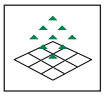

Control of the Physical World by Intelligent Systems

An intelligent agent, interacting with the physical world, must cope with a wide range of demands. Different scientific and engineering disciplines, with different abstractions of the world, have found different "pieces of the puzzle" for the problem of how the agent can successfully control its world. These disciplines include:

- AI (including qualitative reasoning, planning, machine learning, and intelligently guided numerical simulation)
- control theory
- dynamical systems
- fault diagnosis
- fuzzy logic and systems
- neural nets
- computer vision
- robotics

The goal of this symposium is to attempt to understand the puzzle as a whole, by bringing together researchers with experience assembling two or more pieces. The emphasis will be on successful projects in this area that exploit results or methods from several disciplines.

Abstractly, the important questions will be:

- What are the strengths and weaknesses of each piece of the puzzle?
- How do we put together two pieces to exploit their strengths and avoid their weaknesses?
- How do we reconcile the different conceptual frameworks to make different approaches mutually comprehensible?

In order to make our discussions mutually comprehensible, participants should relate their work to one of a small number of everyday tasks: vacuuming the floors in an ordinary house, coping with furniture, pets, trash, etc.; controlling a process such as a pressure-cooker, including set-up, start-up, normal operation, anticipating and handling emergencies, shut-down, and clean-up; automated driving of a car through city and/or highway traffic, including learning spatial structure and using maps; or learning to understand and control one's own sensory-motor system, including seeing, grabbing, walking, running, bicycling, juggling, etc.

The symposium will be organized around a few presentations and lots of discussion. In some sessions, a successful project will be presented and critiqued. In others, a problem will be posed, and relevant contributions collected and evaluated.

The working papers will be distributed in advance, so participants can familiarize themselves with each others' positions before the symposium. We expect conversations of the form: "What problem are you working on?" "Why is that important?" "How can I help you?" and "How can you help me?"

Attendance at the symposium will be limited. Some attention will be given to balance among areas, but the primary criteria will be successful synthesis of multiple approaches to intelligent agenthood, and ability of the participant to communicate across discipline boundaries.

Papers should focus on one of the above everyday tasks (or a task of

similar familiarity and concreteness). It would be helpful to include a glossary of key concepts to help bring the reader into your conceptual framework.

Submission Information

Five copies of either full papers (twenty pages max) or short position papers (five pages max) should be sent to:

Benjamin Kuipers
Cochair, AAAI Intelligent Agent Symposium
Computer Sciences Department
University of Texas at Austin
Austin, Texas 78712 USA

Organizing Committee

Piero Bonnisone, General Electric; Jim Hendler, University of Maryland; Michael Jordan, MIT; Benjamin Kuipers (cochair), University of Texas; Lyle Ungar (cochair), University of Pennsylvania.

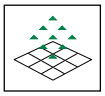

Knowledge Representation for Natural Language Processing in Implemented Systems

This symposium is intended to be a meeting of researchers actively working on implemented knowledge representation and reasoning (KRR) systems for general natural language processing (NLP) in order to assess the current state of that field. Specific topics of interest include the following:

Expressiveness and generality of the representation language with respect to natural language. For example, coverage of complex object descriptions and treatment of quantification. What are the trade-offs in increasing the expressiveness of the representation language to support natural language?

Inference methods that parallel reasoning in natural language. Natural deduction systems, for example, are so called because of the apparent naturalness of the proof procedure. Another example is surface reasoning, based on the syntactic structure of the natural language.

Ability of the formalism or system to capture important semantic and pragmatic aspects of natural language. For example, the computational relationship between the representation language and the parser/generator. Is it possible to define the representation language to facilitate this relationship? At what cost?

How many or kinds of representation languages are needed for general NLP? Many NLP systems actually use two representation languages: a semantic representation language that captures the semantics of a sentence, and a knowledge representation language that is used to do reasoning and represent the system's general knowledge about the domain. Typically these languages are quite different, with the former being a much more powerful language (including modalities, lambda expressions and other higher order constructs, generalized quantifiers, etc) and the latter being what the system's reasoning engine actually operates on, usually something more or less equivalent to first-order logic, (or an even more restricted vivid representation like a relational database) This raises the question of how these languages relate to each other, and if it is possible (or desirable) to have a semantic representation language that supports inference.

Stands on issues such as:

- What problems are solved, and how to use the solution(s).
- What areas need work.
- Defense or attacks of the standard design of morphology-syntax-semantics-pragmatics.

Submissions to the symposium should address these topics by showing some text that an actual implemented system can understand, how the information contained in that text is represented, what background information is used by the system, how that information is represented, how the system processes the knowledge to do interesting things (such as answering interesting questions about the text), and how the information is processed into answers. Reports on projects whose purpose is to simulate human understanding of texts will be preferred over projects

whose purpose is to provide natural language interfaces to databases, planners, or to pragmatically oriented knowledge bases.

The format of this symposium will be designed to encourage interaction amongst the participants. To this end, new, previously unpublished work-in-progress on the topics of the symposium is most desirable, as are stands on the issues outlined above.

Submission Information

Potential attendees should submit an extended abstract of no more than 10 pages (exclusive of references), twelve point, double-spaced, with one inch margins. Submissions not conforming to these guidelines will not be reviewed.

The symposium format will also include one or more panel discussions on the issues listed above. Attendees wishing to participate in panels, or wishing to suggest other panel topics, should indicate their interest and include a current vita along with their extended abstract.

Demonstrations of working NLP/KRR systems are also of interest, however attendees must provide their own hardware and software support. Attendees interested in this option should indicate what they are planning on demonstrating, and how they propose to do so. This information should be provided with their extended abstract in a cover letter that clearly states their interest in this option.

Email submission is preferred, and should be directed to the Symposium Chair at

ssa231f@csm560.smsu.edu and
syali@cs.buffalo.edu

Preferred email submission formats are: stand-alone LaTeX, PostScript, or plain text (for abstracts without complex figures, etc).

If email submission is not possible, then five copies of the paper should be mailed to the Symposium Chair:

Syed S. Ali
Chair, AAAI Fall Symposium on
Knowledge Representation for
Natural Language Processing in
Implemented Systems
Department of Computer Science
Southwest Missouri State University
901 South National Avenue
Springfield, MO 65804
(417) 836-5773

Organizing Committee

Syed S. Ali (chair), Southwest Missouri State University, ssa231f@csm560.smsu.edu or syali@cs.buffalo.edu; Douglas Appelt, SRI International; Lucja Iwanska, Wayne State University; Lenhart Schubert, University of Rochester; Stuart C. Shapiro, State University of New York at Buffalo.

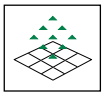

Improving Instruction of Introductory Artificial Intelligence

Introductory artificial intelligence is a notoriously difficult course to teach well. The two most straightforward strategies are to either present a smorgasbord of topics or to focus on one or two central approaches. The first option often yields disjointed and superficial results, whereas the second presents an incorrectly biased picture. Often students come away with a feeling that the subject matter is not coherent and that the field has few significant achievements.

Other issues instructors must face are: What should the balance be between the emphasis on cognitive modeling vs. engineering solutions to hard problems? How should one handle the well-known phenomenon of solutions that work no longer being considered to be part of AI?

What strategies do instructors take to deal with these issues? What are some underlying themes that can be used to help structure the material? Are there sets of principles that can be used to instruct the material, even if they do not precisely reflect all of the current viewpoints?

We propose that the AI community meet and tackle these issues together. The goal of the symposium is to provide an opportunity to discuss these difficult questions, as well as to share successful strategies, problem assignments, instructional programs, and instructional "bloopers".

The symposium will be organized as a workshop, although there will be tutorial aspects, allowing participants to learn from the experiences of colleagues who have worked out successful solutions. All attendees will participate in less formal breakout sessions as well as in discussion following presentations.

We are soliciting four kinds of contributions, corresponding to four presentation types.

Type I

A discussion of a successful strategy for instructing introductory AI, based on experience. The descriptions should be centered around a syllabus and must describe how the strategy handles the smorgasbord versus bias problem, i.e., what underlying unifying themes give cohesiveness to the approach while at the same time covering the material well. Accepted descriptions (one to three pages plus the syllabus) will appear in the working notes, and the authors of some of these will be asked to present these ideas in the form of a talk.

Those submitting papers on successful overall approaches should also address the following issues, and optionally those listed under Type II as well.

1) What are the basics that must be covered, and how are they integrated into the theme of the course?

2) One strategy is to describe a problem to be solved and then discuss several different techniques that attempt to solve the problem. The inverse strategy describes a technique and then explores how well it performs on different problem types. Which, if either, strategy is used?

3) What kind of curriculum does this fit in to? (E.g., isolated undergraduate semester, part of an intelligent systems series, overview for graduate students, etc.)

Type II

A position paper addressing at least one of the following issues (one to four pages). Selected position papers will appear in the working notes and some authors will be asked to participate in well-structured, strongly moderated panel discussions (i.e., answers to a small set of questions prepared in advance).

Content issues include:

1) What is the role of cognitive motivation, if any? Should there be an emphasis on simulating intelligence, modeling the mind, etc.?

2) What is the role of formalism?

3) Should commercially feasible aspects be addressed, and if so, how?

4) How should we address the question, is it still AI if it can be done?

5) What role should advanced areas (e.g., vision, robotics) play?

6) What kind of hooks should be left for more in-depth courses (e.g., graduate ML, NLP, vision, KR, connectionism, etc.)?

7) What is the role of historical developments? Should this be a structuring theme?

Curricular issues include:

1) What kind of department; interaction with other subareas, e.g.: cognitive systems, intelligent systems.

2) What kind of educational program: One semester or quarter overview? A two or three semester series centered around a theme (e.g., KR)?

3) Undergraduate versus graduate: How much should they overlap? Remedial introduction for those who never took undergraduate AI? What are the positive and negative results? Undergrad pedagogical, hands on experience? Graduate reading: primary sources?

Teaching college perspective issues include:

1) What are the special issues?

2) How can pooled resources help?

Laboratory involvement issues include:

1) Is programming useful or a time waster? Is it better to use and observe existing tools instead? Or is a compromise—modifying existing programs—best?

Innovative ideas?

Type III

Informative descriptions of programming tools for assignments and as instructional aids. Many potential symposium participants have expressed interest in acquiring a collection of useful programs both for demonstration of AI concepts and for use in homework assignments in lieu of requiring students to spend their time programming. This is an opportunity to advertise tools, videotapes, and other teaching aids. We would like all participants that have used tools that they have found to be useful to list these tools, and all developers of new tools to describe them. For those who have created new tools, we would like to create a repository of same.

Type III submissions should address at least one of the following (one to five pages):

1) Descriptions of new educational tools specifically designed for instruction of AI (e.g., The FLAIR project of Papalaskari et al.).

2) Recommendations and pans of existing tools as instructional aids.

3) Proposals for tools or tool con-

struction methodologies that don't exist but should.

4) Tools for empirical exploration (e.g., a tutorial guide to the tools in the machine learning repository, as applied for instructional purposes).

Type IV

Contributory questions (one to two pages). Those wishing to participate but who do not wish to contribute in one of the previous categories must indicate their interest by describing what issues they are interested in (including those listed above) and suggesting questions to be addressed by the symposium. Descriptions of bloopers—mistakes made that others should be warned away from, are also appropriate. We are especially interested in topics to be discussed by small groups, and questions to be asked of the panelists.

Submission Information:

Authors may send contributions of more than one type, if desired. Only one document should be submitted. A submitter who wants to make more than one type of submission should simply label each part of the document with the appropriate type (Type I, Type II, etc).

Email submissions are strongly preferred. Send one copy, plain text, to: marti@cs.berkeley.edu. PostScript is acceptable if it is accompanied by a plain text version as well; this way diagrams can be included in an electronic form. Those people who cannot send an electronic version should send three hard copies to:

Marti Hearst
Chair, Symposium on Improving Instruction of Introductory AI
Xerox PARC
3333 Coyote Hill Rd
Palo Alto, CA 94304

Organizing Committee

Marti Hearst (chair), UC Berkeley; Haym Hirsh, Rutgers; Dan Huttenlocher, Cornel; Nils Nilsson, Stanford; Bonnie Webber, University of Pennsylvania; Patrick Winston, MIT.

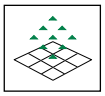

Planning and Learning: On to Real Applications

Planning and learning research have been progressing in parallel over the past several years, but very few research projects have bridged the two areas. However, there is a great deal of benefit from their interaction, especially when they are concerned with real applications. As the complexity of planning problems increases, it becomes of particular interest to identify learning opportunities in order to automate the acquisition of a planner's knowledge in new applications. At the same time, planning problems are a useful testbed and a source of challenges for learning research. The goal of this symposium is to discuss the implications of practical planning applications on both learning and planning research.

The symposium will highlight empirical work on practical problems as an invaluable source for understanding the complexity of the planning task. We expect the analysis and discussion of practical domains to be a solid basis for turning formal planning and learning algorithms into efficient practical ones. As a desirable side effect of the symposium, we also envision the emergence of an initial comparative insight into different planning and learning algorithms from a practical standpoint. In particular we would like to discuss characterizations of application domains, comparisons of performance of different planners on the same task, and practical limitations or power of an approach.

Specific topics of interest for the symposium include:

Practical problems: What makes practical real problems different from simplified simulated tasks? What are the implications of practical problems for specific planning algorithms? What are the learning opportunities for specific planning algorithms to handle practical problems efficiently?

Learning and knowledge acquisition: What learning algorithms were developed or extended to address needs of the application? What tools help extend and maintain planning knowledge?

Learning, planning efficiency, and plan quality: What are learning opportunities for a planning algorithm? How can a planner improve its efficiency based on its past experience? What are measures of plan quality in real applications? How can a planner improve the quality of the solutions it generates?

Scaling up: How well does the approach behave in tasks and problems of increasing size and complexity? What issues need to be addressed and what extensions to the framework are demanded by particular applications?

Domain features: What features of a practical domain stretch the representation language of a planner? What dimensions can be used to characterize and compare application domains? How can search spaces be characterized in terms of the application?

We encourage submissions and

participation of theoretical planning researchers interested in understanding more practical planning problems and the impact of learning in their algorithms, as well as contributions on practical planning applications that shed light on the challenging issues for learning, knowledge acquisition, and representation in planning domains. We especially welcome position papers that present discussions of issues relevant to the symposium as well as those written by teams with complementary research interests.

The symposium will consist of presentations, invited talks, and discussion sessions. In order to encourage participation in the discussions, the organizing committee will put together a list of issues of concern from the submissions, and distribute it to the participants in advance.

Submission Information

To participate, please submit an extended abstract (up to five pages). Authors of accepted abstracts will be invited (but not required) to submit a longer paper for publication in the working notes. Those interested in attending should submit a one- to two-page research statement and a list of relevant publications. Please include your email address in all submissions. Submit four hard copies to:

Yolanda Gil
Cochair, AAAI Fall Symposium on Planning and Learning
Information Sciences Institute
University of Southern California
4676 Admiralty Way
Marina del Rey, CA 90292
(310) 822-1511
(310) 823-6714 (fax)
gil@isi.edu

Organizing Committee

Steve Chien, Jet Propulsion Laboratory; Yolanda Gil (cochair), USC/Information Sciences Institute, gil@isi.edu; Drew McDermott, Yale University; Dana Nau, University of Maryland; Manuela Veloso (cochair), Carnegie Mellon University, veloso@cs.cmu.edu.

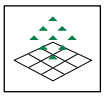

Relevance

Essentially all reasoning and learning systems require a corpus of information to reach appropriate conclusions. For example, deductive and abductive systems use an initial theory (possibly encoded as predicate calculus statements, a Bayesian network or a neural net) and perhaps a to-be-explained observation, and inductive systems typically use both a background theory and a set of labeled samples. With too little information, of course, these systems cannot work effectively. Surprisingly, too much information can also cause the performance of these systems to degrade, in terms of both accuracy and efficiency. It is therefore important to determine what information must be preserved, or more generally, to determine how best to cope with superfluous information. The goal of this symposium is a better understanding of this topic, relevance, with a focus on techniques for improving a system's performance (along some dimension) by ignoring or de-emphasizing irrelevant and superfluous information. These techniques will clearly be of increasing importance as knowledge bases become more comprehensive and real-world applications are scaled up.

There are many forms of irrelevancy. In many contexts (including both deduction and induction), the initial theory may include more information than the task requires. Here, the system may perform more effectively if certain irrelevant facts (or nodes in a neural net or Bayesian network) are ignored or deleted. This realization is clearly related to the notions of abstraction, approximation, qualitative physics, and to the processes underlying many analogical, Explanation-Based Learning, and Case-Based Reasoning systems. In addition, it is often effective to disregard certain links or even nodes in a Bayesian network, after determining that they will have a negligible effect on the distribution of values of a particular variable; similarly one can delete a unit from a neural net if the strengths of its output links are near 0.

In the context of learning, certain attributes of each individual sample may be irrelevant in that they will play essentially no role in the eventual classification or clustering. Also, the learner may choose to view certain samples to be irrelevant, knowing that they contain essentially no new information. This is clearly related to the recent weak learning conversions that provide ways of filtering out some examples and to overfitting.

Yet another flavor of irrelevance arises during the course of a general computation: A computing process can ignore certain intermediate results, once it has established that they will not contribute to the eventual answer; consider alpha-beta pruning or conspiracy numbers in game-playing and other contexts, or control heuristics in derivation.

Submission Information

Potential attendees should submit a one-page summary of their relevant research, together with a set of their relevant papers (pun unavoidable). People wishing to present material should also submit a 2000 word abstract. We invite papers that deal with any aspect of this topic, including characterizations of irrelevancies, ways of coping with superfluous information, ways of detecting irrelevancies and focusing on relevant information, and so forth; and are particularly interested in studies that suggest ways to improve the efficiency or accuracy of reasoning systems (including question-answerers, planners, diagnosticians, and so forth) or to improve the accuracy, sample complexity, or computational or space requirement of learning processes. We encourage empirical studies and cognitive theories, as well as theoretical results.

We prefer plain-text, stand-alone LaTeX or Postscript submissions sent by electronic mail to:

greiner@learning.scr.siemens.com

Otherwise, please mail three copies to:

Russell Greiner
Cochair, AAAI Fall Symposium on Relevance
Siemens Corporate Research, Inc
755 College Road East
Princeton, NJ 08540-6632

Organizing Committee:

Russ Greiner (cochair), Siemens Corporate Research, greiner@learning.scr.siemens.com; Yann Le Cun, AT&T Bell Laboratories; Nick Littlestone, NEC Research Institute; David McAllester, MIT; Judea Pearl, UCLA; Bart Selman, AT&T Bell Laboratories; Devika Subramanian (cochair), Cornell, devika@cs.cornell.edu.