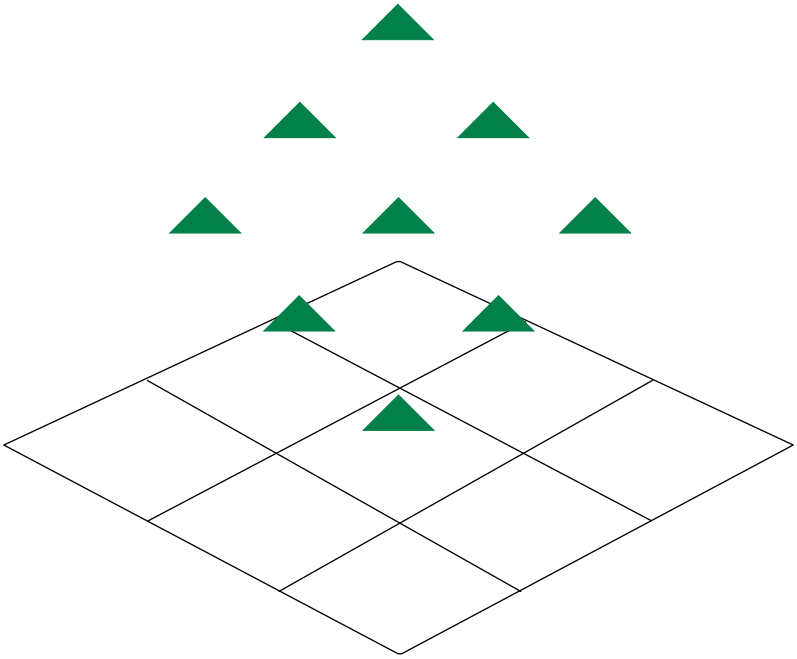

AAAI Fall Symposium Series

CALL FOR PARTICIPATION

October 23, 24, & 25, 1992

Royal Sonesta Hotel

Cambridge, Massachusetts

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

Introduction

The American Association for Artificial Intelligence presents the 1992 Fall Symposium Series, to be held Friday through Sunday, October 23–25, 1992 at the Royal Sonesta, Cambridge, Massachusetts.

The topics of the five symposia in the 1992 Spring Symposium Series are:

- Applications of AI to Real-World Autonomous Mobile Robots
- Design from Physical Principles
- Intelligent Scientific Computation
- Issues in Description Logics: Users Meet Developers
- Probabilistic Approaches to Natural Language

Most symposia will be limited to approximately 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 will be scheduled in which the highlights of each symposium will be presented and an informal reception will be held on Friday evening, October 23.

In addition to invited participants, a limited number of other interested parties will be allowed to register in each symposium. Registration information will be available in July 1992. To obtain registration information write to the address listed on the front of this brochure.

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 May 11, 1992. Acceptances will be mailed by June 8, 1992. Material for inclusion in the working notes of the symposia will be required by August 10, 1992.

Applications of Artificial Intelligence to Real-World Autonomous Mobile Robots

Building intelligent robots that can perform tasks autonomously in the real world has long been a goal of AI research. While a great deal of progress has been made towards this goal in recent years, there is still a wide gap between what theorists claim should be possible and what has actually been demonstrated on real robots.

Various conjectures have been made as to why this gap between theory and practice exists. Some have stated that the theories have simply outpaced the hardware available, but that a few years of engineering effort should suffice to bridge the gap. On the other hand, several researchers engaged in the experimental investigation of real robots have concluded that current theory is based on fundamentally flawed assumptions and cannot be extended to real applications without major changes.

The purpose of this symposium is to bring together AI and robotics researchers to explore the following issues:

- What recent results from the AI research community are ready to be applied to real-world robots?
- For those results which are not yet suitable for implementation, what are the outstanding issues that must be resolved? Are current theories really based on unrealistic assumptions, or are there simpler ways around the difficulties that have been overlooked?
- To what degree can formal theories be applied to practical problems of autonomous robots? Are there limits to what we can characterize mathematically? Can a theory be principled without being formal, and if so, what are the qualities one should strive for in such a theory?

This symposium solicits submissions in the following major categories.

- *Descriptions of formal theories which have been implemented on real robots.*

Submissions in this category should present experimental data and draw conclusions about the applicability of the theory based on this data.

- *Descriptions of theories that aspire to be implemented on real robots but which have not yet been implemented.*

Submissions in this category should make strong arguments as to why the theory will be implementable, and what outstanding issues must be resolved before such implementation is possible. Research performed using simulated or stylized environments should provide strong support concerning their eventual applicability to physical robots.

- *Descriptions of work performed on real robots that does not have a sound theoretical basis.*

Submissions in this category should describe robots performing complex tasks beyond the capabilities of currently implemented theory. Such papers should comment on whether or not it might be possible to give the work a sound theoretical foundation, and what work needs to be done to achieve this.

Results need not be positive; descriptions of outstanding problems—both theoretical and practical—and why current technology is not adequate to solve them are explicitly encouraged. Possible topics for papers include but are not limited to:

- Sensor systems and algorithms
- Navigation algorithms
- Manipulation algorithms
- Dealing with noisy sensor data
- Detecting and recovering from failures
- How to allocate scarce computational resources
- Time-critical planning and control
- Planning in partially unpredictable environments
- Potential applications of autonomous mobile robots
- Case studies of successful and unsuccessful applications

Controversial or informal submissions are encouraged.

Send submissions to:

Marc Slack

The Mitre Corporation, mail stop z401

7525 Colshire Drive

McLean, Virginia, 22102-3481.

Electronic submissions are acceptable only as flat ascii files with no figures. Send electronic submissions to Marc Slack at slack@starbase.mitre.org. All submissions must arrive by May 11, 1992.

Program Committee: Avi Kak, Steve Chien, Marc Slack (cochair), Erann Gat (gat@robotics.jpl.nasa.gov, cochair)

Design from Physical Principles

This symposium concentrates on the use and representation of skills of mathematics, physics and engineering towards the design of all kinds of physical artifacts and processes. The objective of the symposium is to bring together researchers from a diverse set of areas with a common interest in design from physical principles. These areas include, among others, modeling, dynamics, qualitative, temporal, geometric and terminological reasoning, as well as planning, diagnosis, learning, automated deduction, and traditional engineering design. Design provides these researchers a common focus to communicate their ideas, to combine their techniques, and to evaluate their progress.

For the purpose of this symposium, the term *design* is used in a broad, interdisciplinary sense, which includes not only the creation of physical artifacts, but, for example, the creation of discrete and continuous processes, as studied in the planning and qualitative reasoning communities, and the repair or evolution of these creations, as studied within the diagnosis community. Examples include the design of self-repairing devices, chemical processes, controllers of manufacturing processes, and other tasks which draw upon a diverse set of representational and reasoning formalisms.

The symposium focus is on using physical principles—-whether they be scientific, mathematical, or commonsense. Research in the AI design community has largely focused on configuration design, drawing from results on constraint satisfaction and hierarchical planning. Most of this work has not emphasized the physical knowledge, mathematical skills, analytical tools and

commonsense reasoning that comprise a significant component of a designer's abilities. Conversely, areas of AI that deal with core representation and reasoning, such as commonsense reasoning, qualitative reasoning, model-based reasoning, learning and knowledge representation, sometimes lack focus due to their detachment from particular tasks. These areas, together with others such as geometric reasoning, temporal reasoning, dynamics, diagnosis and planning can substantially contribute to the task of design, as they have made substantial progress towards formalizing the use of physical principles.

The symposium aims at creating the appropriate synergy by providing a focused forum for cross-communication. The design community will be served by exposing it to the sophisticated set of reasoning techniques and mathematical tools currently available. The core reasoning communities will be served by providing them with a task which allows them to focus and evaluate their efforts, as well as to provide them with an infusion of new problems, representations and reasoning skills related to design.

All interested participants should submit four copies of a one- to four-page abstract of research interest, focusing on contributions to the topic of the workshop. In addition, researchers wishing to make a presentation should submit four copies of a technical paper of up to about eleven pages, describing completed research or work in progress. We will not insist on new, previously unpublished research, although there will be a preference towards having new material presented. Submissions should be sent no later than May 11, 1992 to:

Brian Williams
XEROX Palo Alto Research Center
3333 Coyote Hill Road
Palo Alto, CA 94304
bwilliams@parc.xerox.com
FAX: 415-812-4334

Program Committee: Brian Williams, Xerox Parc (cochair); Leo Joskowicz, IBM Research (cochair); Tom Dean, Brown University; Jon Cagan, Carnegie Mellon University.

Intelligent Scientific Computation

The purpose of this symposium is to identify the scope of contributions that AI has and might make to scientific computing and vice versa. By *scientific computing*, we mean the practical application of computer programs to solve difficult real-world problems such as the super-computer-sized grand challenges (e.g., forecasting severe weather events, predicting superconductors, or energy conservation and turbulent combustion, as proposed in the Office of Science and Technology report "Grand Challenges: High Performance Computing and Communications"). These problems often involve the formulation and use of mathematical models in the form of differential equations. We include also the significant improvement of solution techniques or solution time for engineering problems routinely tackled today, such as seismic modeling or ship design. We use the term *scientific computation* broadly to include engineering computation as well. Although purely theoretical work on small, highly distilled toy problems is interesting and useful in its own right, we exclude such work from the scope of this symposium unless it also is applied to solving hard engineering and scientific problems.

The result of the symposium should, at a minimum, (1) familiarize AI researchers with the types of problems faced in scientific computing and the work to date addressing the problems by those outside as well as in AI; and (2) familiarize those outside AI with the contributions, past and potential, from AI. In addition, we would like to explore ideas about what further leverage could be gained in future research by specifying well defined interfaces between (quasi) independent components of solution tools, by documenting or standardizing knowledge bases in enough detail to reuse them or reimplement them quickly, and by working with the targeted user community early enough to understand the problems and gain acceptance for the approaches.

Appropriate topics include descriptions of research, completed or in progress, on the application of AI techniques to the various

phases of scientific and engineering computing (e.g., problem definition, solution formulation, software construction, experimentation, analysis, or revision). Equally important are summaries of work from technologies other than AI that impact on the role of AI in scientific computing.

We encourage the submission of summaries of a related technology such as symbolic manipulation systems, compiler optimization, vectorizing/parallelization tools, scientific databases, shared knowledge representations, or numerical libraries. Such a presentation could summarize the current state of the art, the strengths and weaknesses of the approach, and speculate on how that technology could help, replace, or benefit from related AI approaches. Alternatively, a submission might describe contributions of scientific computing to AI, such as a good application problem to work on, or an alternative model of intelligent behavior.

Submissions on intelligent scientific computing should address aspects such as:

- Specifying specific instances of problems to be solved (e.g., by equations, geometries, construction from primitive components, formal languages, domain-specific languages)
- Solution techniques (applying numerical approximations, symbolic manipulation, problem reformulation or specialization, program synthesis, special-purpose simulation generators, adaptive control of the simulation/analysis, knowledge-based methods)
- Representing knowledge of physics (directed at scientific computing tools)
- Software issues such as performance and reliability, modifying models and algorithms, recording design rationales, documentation, and validation
- Visualization or interpretation of results
- Experiment guidance and management (parameter set up, model/data management, computational steering, results maintenance/tracking)

Submissions can describe implemented systems for directly setting up or answering specific questions, simulating classes of problems, or synthesizing programs that solve the problem. Live demos or videotapes may be appropriate for implemented systems.

Submissions of AI approaches should note the competing technologies and should address the specific contribution of the AI approach: Does it enable problems to be set up or solved more efficiently? Does it enable the solution of more difficult problems? Does it interface to or replace existing technologies such as libraries of numerical routines? In considering all submissions describing the use of symbolic techniques, we will use criteria such as the following. Does the submission demonstrate how its techniques significantly facilitate complex numerical computing? What phase(s) of the scientific computation process is (are) being addressed? Is the work being done in the context of a realistic science or engineering problem? If a technique claims to replace numeric computing by symbolic reasoning, has the result been demonstrated to be acceptable to a user who would normally use numeric computing? Surveys of the utility of some of these areas might be more useful than details on specific implementations.

The symposium will most likely be organized around half a dozen topics, determined by the submissions, with each session consisting of a few brief presentations followed by a discussant and then general discussion. Some sessions will allow for long discussion periods.

Potential participants should submit a two- or three-page summary of past and current work in the area organized around the aspects described above. Although a list of related publications should be included, the text description should be thematic rather than simply chronological. In addition, submissions should include a one- or two-page description of a critical issue that the author(s) would like to see addressed by the symposium, including a description of any presentations the author(s) would be willing to make. Additional supporting material in the form of copies of previous papers or extended abstracts is optional.

Submit either email (ascii or LaTeX) source or four paper copies to be received by May 11, 1992 to

Elaine Kant
Schlumberger Laboratory for Computer Science
P.O. Box 200015
Austin, TX 78720-0015
Email: kant@slcs.slb.com
Phone: 512-331-3737
Fax: 512-331-3760

Program Committee: Elaine Kant, Schlumberger Laboratory for Computer Science (kant@slcs.slb.com); Richard Keller, NASA Ames Research Center (keller@ptolemy.arc.nasa.gov); Stanly Steinberg, University of New Mexico, (stanly@crunch.unm.edu).

Issues in Description Logics: Users Meet Developers

During the 1980s, a branch of the knowledge representation community developed a specialized class of logics known either as terminological logics or description logics (DLs). Numerous implementations of these logics have emerged, along with a substantial body of theoretical work. This workshop is intended to bring together developers of description logic systems with users of these systems. We are looking for issues and ideas that provide new ways of viewing the existing DL paradigm, comment on proper ways to apply the paradigm, or describe representational extensions or alternatives that complement a DL reasoner.

We plan to adopt a workshop environment—during the discussions individuals will be given an opportunity to present their views, but there will be no timetable of presentations. We are soliciting position papers for the following areas:

• *Theory meets practice:* Putting DL theory into practice means building a reasoner (i.e., a classifier and recognizer), and then adding anything else necessary to make the system useful. In some cases (e.g., the semantics of retraction), there seems to be reasonable agreement on how to implement the desired functionality, while for others (e.g., role-closing constructs, managing inconsistency) there is considerable disagreement. We specifically invite discussion on structural subsumption classifier architectures with regard to issues of completeness, characterizability, and performance.

• *Methodological principles:* A general and formal methodology for building knowledge bases using DLs has yet to be developed. Such a methodology should try to address questions such as: What constitutes a good Tbox? Could we define a style checker for a terminological knowledge base? Do constraints belong inside or outside of primitive definitions? Is there a difference between a role and a binary predicate, and if so, what is it? Can we recapture what was lost when we moved from frame representations to terminological concepts and descriptions?

• *Representational extensions:* This topic includes extensions that make a DL more expressive or make it simpler to say things in the logic, but that stay within or close to the boundaries of standard logic. Possible issues for discussion include generalized or alternative notions of inheritance (e.g., part-of inheritance), reification, second-order (meta-level) reasoning, temporal reasoning, and representations of specialized classes of objects (e.g., sets and sequences). In all cases, the extension should not violate the essential character of a DL. However, identifying this essential character could itself be a topic of discussion.

• Integration with other reasoning paradigms: Under this topic we consider the juxtaposition of the DL paradigm with other non-logical or non-classical forms of reasoning. We are specifically interested in systems that interface a DL classifier and some form of approximate reasoner (note: merely *extending* a classifier to perform plausible reasoning is less interesting). Another appropriate topic is the integration of a DL classifier with a planning system.

We are specifically looking for position papers that express *opinions* on the right way to design some aspect(s) of a DL-based KR language or system. Authors should include a brief rationale supporting each of their positions. We are as much interested in the why as in the what. Statements such as “I am interested in talking about X” or “My lan-

guage looks like this" are *not* what we are looking for. For each session, we hope to identify a pair of speakers with semi-opposing viewpoints to make introductory statements at the start of the session. The progress of the remainder of a session will be guided by a session moderator. Attendees will be given a chance to present a few slides backing their particular viewpoint at the time it becomes relevant within the discussion.

The recommended length for papers is between 500 and 2000 words. Each paper should have a single author. We will circulate accepted electronic versions to all participants prior to the workshop. Electronic submissions are preferred, and should be sent to macgregor@isi.edu by May 11, 1992. If electronic mail is unavailable, four paper copies should be submitted to:

Robert MacGregor
USC/Information Sciences Institute
4676 Admiralty Way
Marina del Rey, CA 90292-6695

Program Committee: Robert MacGregor (chair), Deborah McGuinness, Eric Mays, Tom Russ.

Probabilistic Approaches to Natural Language

Recently there has been a resurgence of interest in probabilistic methods in AI, spurred by technical developments which have made these methods more practical. Bayesian and decision-theoretic approaches have been facilitated by the development of graphical representations such as belief (or Bayesian) networks, and influence diagrams. Learning approaches have been promoted by new developments in statistical learning (particularly hidden Markov models). These methods all offer hopes to address problems of brittleness and knowledge representation in natural language processing. Each has its own special strengths, however. Bayesian approaches have a clear conceptual framework and powerful representations, but must still be knowledge-engineered, rather than trained. Hidden Markov models have a clear conceptual framework and the ability to learn, but structure must be given, and the model is weak.

This symposium will bring together researchers applying both of these probabilistic methods in order to share perspectives. We intend that the discussion will emphasize reviews of the current state of the art and views of the most promising lines of research.

Of particular interest are novel applications of statistical and Bayesian techniques, systems which add more complicated knowledge representations to statistical methods or adaptive Bayesian methods. We are also interested in research where Bayesian and statistical methods are use to solve foundational issues in knowledge representation, natural language semantics and acquisition of semantic representations.

Some examples of such research are:

- The use of statistical methods to extract information from text.
- Combining primary source evidence from large corpora with dictionary knowledge for various applications, including part of speech tagging, sense discrimination/disambiguation, and bilingual word/phrase matching.
- Using Bayesian methods to implement abductive approaches to NL interpretation.
- Probabilistic approaches to machine translation.
- Combining high-level knowledge with low-level speech recognition.
- Indexing and retrieval of concepts from text.

Those wishing to attend the symposium should submit a one-page statement of research interests and accomplishments, and a bibliography of selected publications. Those wishing to present their work for discussion should submit, in addition, an extended abstract of no more than four pages.

Potential participants should submit these materials by electronic mail to rpg@cs.tulane.edu. If this is for some reason impossible, four copies of a printed docu-

ment may be submitted to:

Robert Goldman
Computer Science Department
301 Stanley Thomas Hall
Tulane University
New Orleans, LA 70118-5698 U.S.A.

All submissions must be received by May 11, 1992.

Program Committee: Robert Goldman (chair), Peter Norvig, Eugene Charniak, Bill Gale.