

AAAI 1992 Fall Symposium Series Reports

■ The American Association for Artificial Intelligence held its 1992 Fall Symposium Series on October 23–25 at the Royal Sonesta Hotel in Cambridge, Massachusetts. This article contains summaries of the five symposia that were conducted: Applications of AI to Real-World Autonomous Mobile Robots, Design from Physical Principles, Intelligent Scientific Computation, Issues in Description Logics: Users Meet Developers, and Probabilistic Approaches to Natural Language.

The American Association for Artificial Intelligence (AAAI) held its 1992 Fall Symposium Series on October 23–25 at the Royal Sonesta Hotel in Cambridge, Massachusetts. This article contains summaries of the five symposia that were conducted: Applications of AI to Real-World Autonomous Mobile Robots, Design from Physical Principles, Intelligent Scientific Computation, Issues in Description Logics: Users Meet Developers, and Probabilistic Approaches to Natural Language. Technical reports of the symposia Applications of AI to Real-World Autonomous Mobile Robots, Intelligent Scientific Computation, and Design from Physical Principles are available from AAAI. Instructions and an order form for purchasing electronic and hard-copy versions are printed in the *AI MagFax*.

Applications of AI to Real-World Autonomous Mobile Robots

This was the second AAAI Fall Symposium focusing on AI and robotics. The focus was on autonomous

mobile robots for unstructured uncertain environments (last year's symposium was on sensory aspects of intelligent robotics). Robotic operation in such environments presents unique challenges for AI because it requires interaction with environments that are difficult or impossible to model accurately and precisely. As before, the symposium generated much excitement and enthusiasm. A great deal has been accomplished in the last five years, and there is a general consensus that additional progress will be made in the near future.

Perhaps the most exciting development in the field is the recent availability of inexpensive off-the-shelf robots that include enough sensors and processing power to do interesting experiments. As a result, a large number of people are beginning to experiment with real robots, which is both good news and bad. The bad news is that a great deal of theoretical work turns out to be based on assumptions that are unrealistic and difficult or impossible to discharge but that retain the utility of the theory. The good news is that getting mobile robots to do interesting and useful things turns out to be surprisingly easy if you follow Nike's advice: Just do it. Mobile robot navigation is, at least in practice and in restricted (though realistic) situations, largely a solved problem. Robots navigating autonomously in office environments are being produced with gratifying regularity. An interesting result of this work is the uniformity of the approaches being used on successful robots. By and large, working robots are being produced using a multilayer approach based on a set of compliant processes

that couple sensors to actuators under the control of a higher-level symbolic reasoning mechanism. A number of essentially identical architectures following this general approach have been developed independently and implemented successfully over the last two years.

Probably the greatest challenge lying ahead for the field is translating these practical successes into general scientific results and principles. Indeed, some participants argued that it was premature to even attempt this translation step. It was argued that robotics should follow the route of the natural sciences, which historically have concerned themselves with studying and explaining natural phenomena. In the case of mobile robots, we must produce the phenomenon before we can study it; so, we must, temporarily at least, abandon our quest for theories and conduct ourselves more like engineers (or architects or alchemists) if we are to make progress. A small but vocal minority argued that theoretical work should not be abandoned because it can be used to predict failures before they occur and, thus, save time and effort.

Finally, a consensus was reached that it is necessary to develop some principled experimental methodology through which ideas can be tested and convincingly confirmed or refuted. Both real robots and simulation should be used: simulations to perform large numbers of experiments economically and physical robots to verify that the simulated results hold up in reality.

The general consensus of the symposium attendees was that the gathering resulted in a successful interchange

of ideas. We want to hold another symposium, at which we hope to address some of the wider concerns of the AI community as they relate to the design of intelligent robots.

Erann Gat

Jet Propulsion Laboratory
California Institute of Technology

Marc Slack

MITRE Corporation

Avi Kak

Purdue University

Steve Chien

Jet Propulsion Laboratory
California Institute of Technology

Design from Physical Principles

Design is a complex human activity that requires a model of the surrounding environment and a way to reason about the behavior of devices embedded within this environment. The Design from Physical Principles Symposium focused on the use and representation of skills from mathematics, physics, and engineering in the design of physical artifacts and processes. The objective of the symposium was to bring together researchers from a diverse set of areas with a common interest in design from physical principles. These areas included modeling, dynamics, qualitative, temporal, geometric, and terminological reasoning; planning; diagnosis; learning; automated deduction; and many aspects of traditional engineering design. Design provided researchers with a common focus to communicate their ideas, combine their techniques, and evaluate their progress.

The symposium was organized around the different stages of the design process, which can roughly be characterized as a generate-evaluate-modify cycle. Given a set of specifications and constraints, generative design produces a candidate design. To evaluate the candidate design, one or more models capturing the properties of interest are generated. Simulation and analysis are then performed to further evaluate the candidate design and identify potential conflicts. Based on the results of

the evaluation, the design is modified or adapted, and the evaluation cycle is repeated. At the symposium, specific approaches addressing one or more stages were proposed. They mainly differed in the type of artifacts designed (plans, lumped-parameter systems, devices) and the representation methods used (discrete events, ODES, PDES, and so on). A separate session was devoted to design in planning and control.

To gain some perspective on the field, we organized two sessions to examine both past and current design research projects. These projects included the Massachusetts Institute of Technology (MIT) designer's apprentice; the Rutgers University Ship Design Project; the Stanford University How Things Work Project; and broad efforts from University of California at Berkeley, Cornell University, and Carnegie Mellon University. Two keynote addresses were presented by G. Sussman and W. J. Mitchell, both from MIT. Sussman's talk emphasized the use of sophisticated qualitative analysis of nonlinear dynamic systems in design. Mitchell showed how shape grammars can be used to capture the form and function of buildings and devices. To give the workshop participants a practical flavor of design from physical principles, we organized a design project based on a realistic problem—the design of actuators for a self-adapting mirror at the Lawrence Livermore National Laboratory. Participants worked on the problem during the workshop in mixed teams of engineers and computer scientists and presented their results at the end.

The workshop was successful in bringing together researchers from various disciplines and having them express their differing perspectives. Although some areas, such as modeling, analysis, and optimization, have made substantial progress, other more encompassing areas, such as design generation and adaptation, are still in their infancies. The most resounding theme of the symposium was the crucial role that realistic design scenarios play in exposing important research problems.

Leo Joskowicz

IBM T.J. Watson Research Center

Brian Williams

XEROX Palo Alto Research Center

Jon Cagan

Carnegie Mellon University

Tom Dean

Brown University

Intelligent Scientific Computation

The style of scientific computing is poised for a change. Traditionally specified in FORTRAN by the practicing scientist or engineer, scientific computations are becoming more complex because parallel architectures enable increasingly ambitious simulations. The application of AI techniques can free mathematical modelers from tedious programming tasks and enable them to carry out more complex simulations and analyses.

The purpose of this first symposium on intelligent scientific computation was to identify the scope of contributions that AI has made and might make to scientific computing and vice versa. The workshop participants' collective background was a good mix of AI and practical application experience, with many individuals having experience in both. The applications discussed at the symposium ranged from fluid dynamics and particle accelerators to yacht design to the analysis of AI programs themselves. These are difficult problems that challenge the AI community to demonstrate its maturity. A wide range of AI techniques have been applied. For example, backward chaining and abstraction were used to formulate mathematical models; expert systems and program synthesis were used to select or generate codes to solve the models numerically. A novel method for accessing trade-offs among selection criteria was presented. Other reasoning techniques included assumption maintenance to track approximations made and their dependence on problem characteristics or efficiency trade-offs, spatial reasoning to generate the grids over which approximations are applied, and qualitative reasoning and learn-

1993 Conference on Artificial Intelligence in Petroleum Exploration and Production

May 19 – 21, 1993

Dallas, Texas

For Conference Information:

(214) 754-6384 (voice)

(214) 754-3016 (fax)

In cooperation with AAAI

ing that help in data analysis. Much scientific knowledge is represented mathematically and procedurally, with the need for a language that combines mathematical and programming constructs. Object-oriented representations and transformation rules are also important representation techniques.

Although the presentations typically focused on successful applications of existing AI technology, the discussion periods addressed problems and possible research directions. For example, although symbolic algebraic systems are powerful tools used by many of the participants, they need to be augmented to include representations for multiple solutions and assumptions made by the system. The importance of

validating solutions was stressed, and a desire for better spatial reasoning techniques was expressed. Although there were a few examples of systems that discovered new scientific results, a certain lack of boldness in ambition was noted. Finally, the workshop participants discussed the possibilities for shared representations of specifications and knowledge bases. The problem of applying AI techniques to scientific computation is too hard to solve without collaboration.

Elaine Kant

Schlumberger Laboratory for Computer Science

Richard Keller

NASA Ames Research Center

Stanly Steinberg

University of New Mexico

Issues in Description Logics: Users Meet Developers

The typical knowledge representation system is like an iceberg—most of it is hidden from view. Its internal architecture is invisible to its users as well as to the AI literature. Thus, the recent symposium on description logic systems provided a valuable setting for bringing together users and developers to discuss their particular brand of knowledge representation system.

A description logic (also known as a KL-ONE-like language or a taxonomic logic) is a subset of first-order logic whose expressivity has been restricted to the point where competent and reasonably efficient deductive inference is possible. Alternatively, we might characterize a description logic as a funny representation language that has no variables. For example, a typical description logic might define the concept of mother as follows:

(defconcept Mother (and Woman
(atleast 1 child))) .

This concept states that a mother is a woman who has at least one child. A description logic system is designed to reason with such descriptions and with term definitions that bind concept names to descriptions. A description classifier organizes a set of descriptions into a taxonomy that is ordered by subsumption relationships computed between descriptions.

Description logics represent an active area of current knowledge representation research; the 1992 Knowledge Representation Conference devoted three sessions to papers on description logics. One reason for this interest is that many description logics have attractive theoretical properties. Perhaps a better reason for interest in description logics is that a description classifier turns out to deliver useful inferences that are needed by real applications. This symposium occurred at a point that might represent a transformation in the evolution of description logic technology. Description logics are now considered to be well understood, and many implementers are beginning to broaden their horizons with respect to the kinds of knowl-

edge representation issues that they are willing to tackle. Here, we highlight some of the issues that were discussed at the workshop.

Part-Whole Relations: Description logic researchers have beaten the is-a link into submission. Their next goal might be to formalize inheritance relations for part-whole links.

Expressivity: In the absence of completeness as a natural boundary on the expressiveness of a description logic, no one was able to identify theoretically based criteria for determining how expressive a description logic should be. However, this lack of criteria no longer appears to be a deterrent for most implementers of incomplete description logic systems.

Scaling Up: Some researchers discussed plans to apply their systems to knowledge bases containing more than 100,000 facts or more than 100,000 concepts. Scaling up to this size will likely require new technological developments.

Default Logic: A number of efforts are now under way to combine a description logic with some form of default logic. Because these efforts represent ways to actually use default logics rather than just talk about using them, we hope to see concrete progress in this area.

Planning: Many planning systems construct and reason with plan taxonomies. Because description logic systems are proficient at computing taxonomic relationships, several researchers have combined the two technologies, using a description classifier to construct a plan taxonomy.

Prepared presentations occupied only a third of each workshop session, with the remaining time reserved for brief position statements and open discussion. The participants agreed that this format worked extremely well—discussions were for the most part lively, and the workshop proved to be a good medium for exposing both current thinking and future directions for description logic systems.

Robert MacGregor
USC/Information Sciences Institute

Eric Mays
IBM T.J. Watson Research Center

Deborah McGuinness
AT&T Bell Laboratories

Tom Russ
USC/Information Sciences Institute

Probabilistic Approaches to Natural Language

The symposium on probabilistic approaches to natural language was concerned with the use of probability theory for knowledge acquisition and uncertainty management in natural language processing, in particular, the acquisition of linguistic knowledge in assembling resources such as lexicons and grammars. Uncertainty management arises in the need to evaluate and choose between different interpretations of texts.

During the symposium, reports were presented on work addressing problems such as tagging words with parts of speech using statistical methods and learning grammars from corpora and word-sense disambiguation (determining intended sense of a particular word in a text). Statistical methods have recently become useful in attacking these problems because of the availability of large online corpora (on the order of millions of words and still rising) and other machine-readable linguistic resources such as dictionaries and thesauri.

Work on the use of subjective (Bayesian) methods for uncertainty management in natural language processing was also discussed. These methods were in the minority, and despite interest in doing so, it was difficult to see how to tie them to statistical methods. To get a sense of the difficulty of the problems that were addressed, Bill Gale pointed out that after 40 years of work on the problem of word-sense disambiguation, researchers still do not have a clear definition of the problem or even a clear sense of what the senses are to choose from. Nevertheless, a number of conclusions can be drawn based on the

work presented in the sessions. Some of the least controversial of these conclusions were the following:

First, simple statistical methods work much better than more intricate, more brittle rule-based systems. For example, consider the problem of tagging words with parts of speech. Just picking the a priori most likely tag for each word, one seems to get about 90-percent accuracy. Bi- and tri-gram models get over 95-percent accuracy, arguably as good as human performance on this task.

Second, at least currently, there is a distinction between language modeling for prediction (for example, for speech recognition) and interpretation. It is not clear whether this distinction will persist.

Third, one must not undervalue the importance of the words themselves when trying to capture facts about language. Such information is often lost when one abstracts to talk about phrase markers. This statement seems to be a partial explanation for the failure of many parsers and the somewhat disappointing results in learning grammars from text. Yves Schabes (Mitsubishi Research Labs) and John Lafferty (IBM T.J. Watson Research Center) presented various methods for “lexicalizing” grammars that hold promise for overcoming this problem.

Conversely, though, there is a real need for appropriate abstractions to allow us to manage the immense parameter spaces needed for stochastic models of language. This need is all the more essential because of Zipf’s Law. Zipf’s Law states that much of the weight of the probability distributions we care about is in the tails, making the number of trials necessary to thoroughly sample prohibitively large, even given the resources that are now becoming available.

Robert P. Goldman
Tulane University
Eugene Charniak
Brown University
William A. Gale
AT&T Bell Laboratories
Peter Norvig
Sun Microsystems Laboratories