

The Workshop on Logic-Based Artificial Intelligence

Jack Minker

The Workshop on Logic-Based Artificial Intelligence (LBAI), sponsored by the National Science Foundation and the American Association for Artificial Intelligence, was held in Washington, D.C., on 13 to 15 June 1999. The workshop was organized by Jack Minker and John McCarthy. The Program Committee members were Krzysztof Apt, John Horty, Sarit Kraus, Vladimir Lifschitz, John McCarthy, Jack Minker, Don Perlis, and Ray Reiter. The purpose of the workshop was to bring together researchers who use logic as a fundamental tool in AI to permit them to review accomplishments, assess future directions, and share their research in LBAI. This article is a summary of the workshop.

The areas selected for discussion at the workshop were abductive and inductive reasoning, applications of theorem proving, commonsense reasoning, computational logic, constraints, logic and high-level robotics, logic and language, logic and planning, logic for agents and actions, logic of causation and action, logic, probability and decision theory, non-monotonic reasoning, theories of belief, and knowledge representation.

Sessions consisted of one to four lectures. Following the lectures in each session, a discussion was held to assess accomplishments and directions of research related to the topic of the session. In addition, 12 different logic-based systems were discussed briefly, and the systems were demonstrated. A brief synopsis of the lectures and the systems demonstrated follows.

Description of Workshop

In "Abductive and Inductive Reasoning," F. Buccafurri, T. Eiter, G. Gottlob, and N. Leone provided a first step toward integrating model-checking methods with AI. They developed an abstract concept of a system repair problem that is solved by applying logic-based abduction. S. H. Muggleton and F. A. Marginean noted that data-refinement operators in inductive logic programming have been viewed as functions from clauses to sets of clauses. Searches carried out by genetic algorithms allow larger changes by com-

The Workshop on Logic-Based Artificial Intelligence (LBAI) was held in Washington, D.C., on 13 to 15 June 1999. The workshop was organized by Jack Minker and John McCarthy. Its purpose was to bring together researchers who use logic as a fundamental tool in AI to permit them to review accomplishments, assess future directions, and share their research in LBAI.

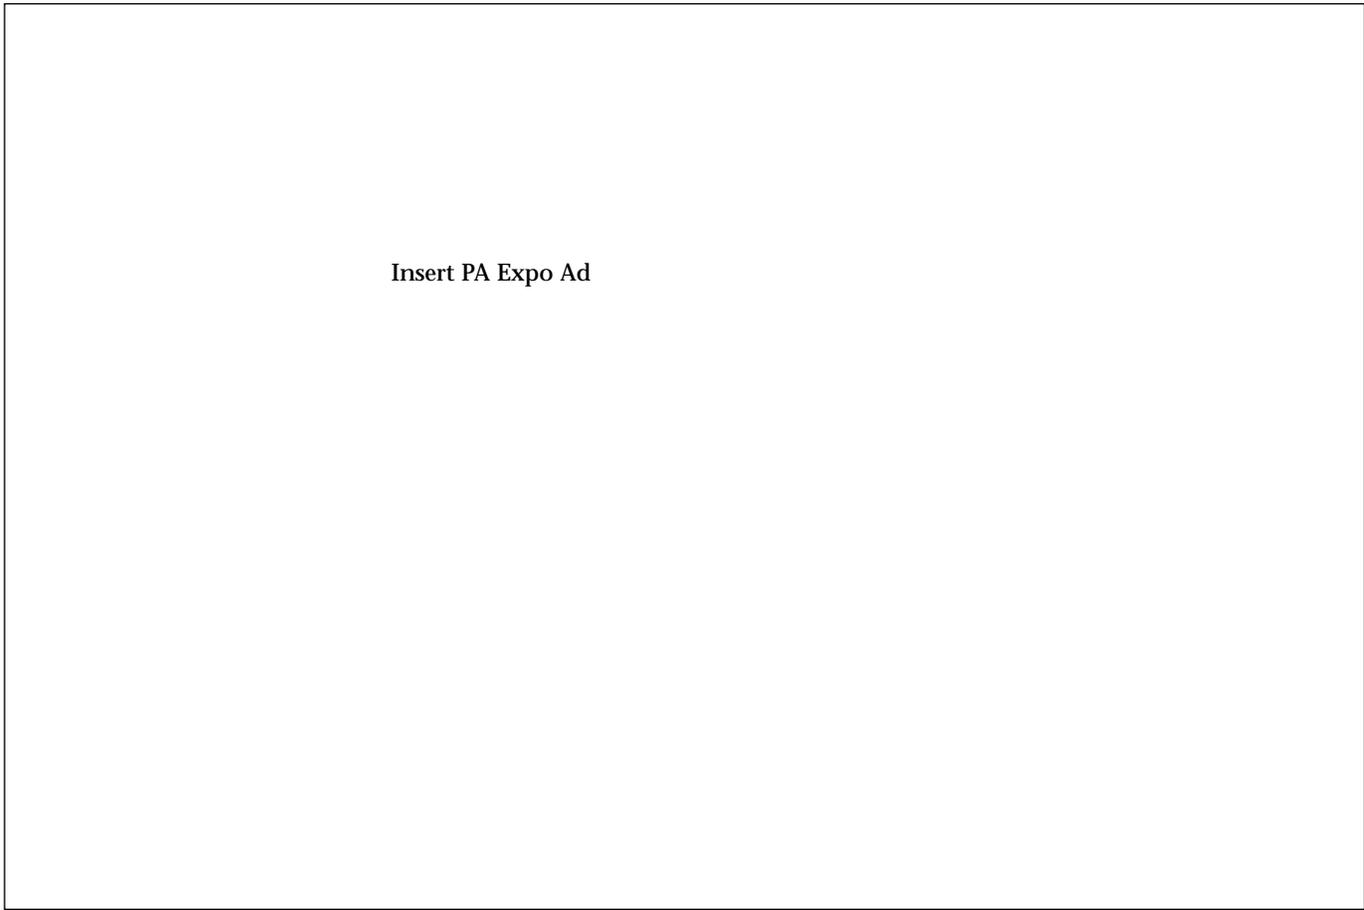
binning pairs of search nodes. They introduce and discuss binary refinement operators within the subsumption lattice that are functions from pairs of clauses to sets of clauses.

In "Applications of Theorem Proving," J. Moore discussed the capabilities of the ACL2 system (the successor of the Boyer-Moore NQTHM prover). ACL2 defines the state of the art in microprocessor verification. It has been used to formally model a Motorola CAP digital signal processor

and the Rockwell-Collins JEM1 microprocessor (the world's first silicon JAVA virtual machine). The use of ACL2 uncovered at least four hardware design bugs—bugs that had survived hundreds of millions of test cases—in time to fix them before the processor was shipped. A. Levy addressed data-integration systems that require a flexible mechanism to describe contents of sources. He surveyed the main logic-based languages that have been considered for describing data sources in data-integration systems and the specialized inference techniques developed in this context. P. Nayak and B. Williams dealt with a reactive self-configuring system installed in the National Aeronautics and Space Administration Deep Space 1 spacecraft now in flight. This autonomous system performs failure analysis and correction for a variety of tasks during the years of active service of the spacecraft. D. Gunning spoke about knowledge base projects at the Defense Advanced Research Projects Agency, where general reasoning systems are needed both to handle the general question-answering problem and to reason from first principles in unanticipated situations.

In "Commonsense Reasoning," McCarthy noted that logical AI involves representing knowledge of an agent's world, its goals, and the current situation by sentences in logic. He characterized a large number of concepts that have arisen in research in logical AI. E. Sandewall addressed what the appropriate research methodology is when logic is used in our field. He argued that any research methodology must specify what is to be counted as relevant results in the research and what is the evidence that is expected to be seen to accept proposed results.

In "Computational Logic," G. Gottlob discussed an application of a restricted form of logic programming admitting negation but not function symbols to hardware verification. The idea is to use a fragment of this restricted form that can be processed in linear time (data complexity). The fragment is relevant to databases because it allows constructs such as transitive closure and, thus, recursive queries and to AI concerns such as frame problems,



Insert PA Expo Ad

inheritance reasoning, and various forms of default reasoning. Apt discussed an application of full first-order logic as a programming language. Although first-order logic is undecidable, by a slightly nonstandard interpretation of logic, one can build a correct (but incomplete) algorithm for example finding (for open theories, the analogous technique is based on the so-called Green trick for finding answer substitutions). A programming language called ALMA-0 was described that combines advantages of both logic (declarativeness) and usual imperative programming. ALMA-0 is based on a three-valued semantics, and consequently, it is correct but not complete.

In "Constraints," R. Dechter described bucket elimination, an algorithmic framework that generalizes dynamic programming for many complex problem-solving and reasoning activities. The bucket-elimination framework was demonstrated for processing knowledge bases such as constraint networks and the Davis-Putnam

method for satisfiability.

In "Knowledge Representation," H. Levesque explored two ways of using expressive logical representation languages without necessarily sacrificing tractability in handling open-world reasoning. Leone discussed how to represent and solve problems in disjunctive logic programming, using the DLV system based on a nondeterministic "guess-and-check" paradigm.

In "Logic for Agents and Actions," J.-J. Meyer applied dynamic logic to view an agent's actions. Meyer also used deontic logic to handle when an agent is permitted to carry out an action. V. S. Subrahmanian applied logical knowledge base techniques and deontic logic. M. Shanahan applied the event calculus to handle agent's actions. In carrying out its actions, an agent must often be able to explain the current situation and assimilate new (sometimes conflicting) information. A variety of techniques are used for this, including abduction (Shanahan), belief revision (Meyer) and logical knowledge base

techniques (Subrahmanian).

In "Logic and Causation," Lifschitz discussed the interaction between the design of languages for describing actions, the analysis of causal reasoning, and the use of propositional problem solvers for planning. He described the development of his work based on these lines of research. H. Geffner focused on the application of theories of action research to planning. He emphasized the need for action languages to facilitate modeling and computation for planning. He illustrated his ideas with the presentation of Functional STRIPS and showed the high impact on efficiency that modeling based on this language has on planning tasks.

In "Logic and High-Level Robotics," C. Baral and M. Gelfond discussed the design of "software components of reasoning, planning, and acting in a changing environment." The dynamic system is viewed as a transition system and described by a logic program. An application to space shuttle control was

described. Reiter provided an overview of the Cognitive Robotics Project at the University of Toronto led by himself and Levesque. The project aims to provide a logical language called GOLOG for specifying and programming the behavior of an autonomous agent.

In "Logic and Language," L. K. Schubert spoke about the situations we reason about and argued for associating situations (including events, episodes, eventualities) with arbitrarily complex sentences rather than just atomic predicates in the process of understanding natural language utterances. He argued that two such notions are needed: (1) a sentence characterizing a situation and (2) a sentence partially describing a situation. He formalized these notions in FOL**, an extension of first-order logic and a part of his episodic logic. R. H. Thomason presented a logic for intra-agent modality, showed how it can be used to formalize reasoning about the attitudes of other agents, and explained how it can be applied to the problem of mutual attitudes.

In "Logic and Planning," H. Kautz and B. Selman discussed BLACKBOX, a planning system that unifies the planning as a satisfiability framework with the plan graph approach to STRIPS planning. B. Nebel discussed the expressive power of preconditions in planning. He showed that preconditions in conjunctive normal form add to the expressive power of propositional STRIPS. He further showed that they do not add any expressive power once there are conditional effects.

In "Logic, Probability, and Decision Theory," it was argued that intelligent agents must be able to deal with uncertainty. H. Prade and D. Dubois developed a logic for making decisions based on qualitative information about preferences and likelihoods (where the likelihoods were modeled using possibility theory). Although the logic was a propositional one, they showed how some of the intuitions underlying decision theory can be embedded in a logical language. C. Boutilier showed how logical ideas and representations can play a role in speeding up probabilistic computations, in particular the computation of policies in Markov decision problems.

In "Nonmonotonic Reasoning," M.



Artificial Intelligence and Mobile Robots: Case Studies of Successful Robot Systems

*Edited by David Kortenkamp, R. Peter Bonasso,
and Robin Murphy*

This book contains thirteen case studies of successful mobile robot systems representing the best available implementations by leading universities and research laboratories. These are not robots that simply work in the laboratory under constrained conditions. They have left the lab and been tested in natural and unknown environments. Many common themes are apparent throughout the book including navigation and mapping, computer vision, and architecture. Each case study is self-contained and describes a complete robot system, including detailed descriptions of important algorithms, and pseudo-code. This book serves as a recipe book for designing successful mobile robot applications. The final case studies in this book describe an approach to using mobile robots in the classroom.

ISBN 0-262-61137-6
\$39.50 softcover

400 pp., illus., bibliography, index
Copublished by The MIT Press

The AAAI Press • Distributed by The MIT Press
Massachusetts Institute of Technology,
5 Cambridge Center, Cambridge, Massachusetts 02142

To order, call toll free: (800) 356-0343 or (617) 625-8569.

MasterCard and VISA accepted.



Seventeenth National Conference on Artificial Intelligence

Austin Convention Center • Austin Texas • July 30 – August 3, 2000

Paper Submissions by January 18, 2000

www.aaai.org/Conferences/National/2000/

Denecker, W. Marek, and M. Truszczyński described an algebraic approach to the semantics of nonmonotonic logics. This work is based on Fitting's theory of logic programs with negation as failure and generalizes this theory to other fix-point formalisms. J. Delgrande and T. Schaub argued that default logic is more general and more widely applicable than previously supposed. They showed how to express preferences among defaults by means of naming default rules and encoding conditions for their invocation.

In "Theory of Beliefs," Perlis gave an overview of what sorts of things could be represented using a formal logic of beliefs and discussed why it was natural to do so. J. Halpern showed that a simple, natural mathematical mechanism served to unify major aspects of belief treatments developed by several theorists, starting from quite different perspectives.

Several system implementations were demonstrated. These systems

were for robotics, deductive databases, planning, inductive inference, and nonmonotonic reasoning. They are as follows:

S MODELS—Developed by I. Niemela, it implements stable semantics for **DATA-LOG** programs with negation.

XSB—Developed by D. S. Warren, it implements full first-order logic programming with tabling under the well-founded semantics.

DLV—Developed by Eiter and Leone, it implements disjunctive logic programming under the answer-set semantics.

DERES—Developed by W. Marek and M. Truszczyński, it implements full default logic of Reiter.

LDL++—Developed by C. Zaniolo, it implements an extension of stratified databases and has additional features such as aggregation for deductive databases.

GOLOG—Developed by Reiter and Levesque, this system focuses on high-level support for cognitive functions

in robotics.

IMPACT—Developed by Subrahmanian, it provides a platform and environment for agent and software interoperability.

CAUSAL CALCULATOR—Developed by N. McCain, it is used for query answering in planning domains using a high-level action-description language. The problem is reduced to satisfiability checking using special-purpose systems.

GENERAL PLANNING TOOL (GPT)—Developed by B. Bonet and H. Geffner, it is a software tool that supports a general approach to planning, including classical, contingent, and probabilistic planning.

TLPLAN—Developed by F. Bacchus, it is a planning system based on model-checking logical formulas against finite models.

BLACKBOX—Developed by Selman and Kautz, it is a planning system that works by converting problems specified by **STRIPS** notation into Boolean satisfiability problems and solving the prob-



The Mexican AI Conference

Acapulco, Mexico
April 10-14, 2000
Hyatt Regency Acapulco

<http://www-cia.mty.itesm.mx/micai2000>

Sponsored by the Mexican Society for Artificial Intelligence and the Acapulco Institute of Technology, in cooperation with AAAI and SMCC and support from IPN, UNAM, ITAM, UDLA, LANIA, UV, CICESE, BUAP, CINVESTAV, ITESM.

This conference is the merging of the Mexican Meeting on Artificial Intelligence (RNIA) and the International Symposium on Artificial Intelligence (ISAI). Both conferences have been held annually since the middle 80's.

INVITED SPEAKERS

There will be 6 keynote speeches by leading AI researchers.

TECHNICAL PROGRAM

Theoretical and application oriented paper presentations.

The conference proceedings will be published in Springer-Verlag's Lecture Notes on AI series.

Topics:

Technologies

Expert Systems and Knowledge Based Systems, Intelligent Agents and Distributed AI, Knowledge Management, Intelligent Interfaces, Computational Intelligence, Learning, Knowledge, Reasoning, Intelligent Robotics, Hybrid Intelligent Systems, Other Technologies.

Application Areas

Business and Commerce, Banking and Finance, Manufacturing and Engineering, Telecommunications, Sustainable Development, Energy and Environment, Education, Construction, Other.

ACCOMMODATION

Special rates for MICA I delegates at the Hyatt Regency Acapulco, Acapulco, Guerrero, Mexico.

<http://www-cia.mty.itesm.mx/micai2000/accomodation.html>

Information about the Hyatt can be found at: <http://www.hyatt.com/pages/a/acaraa.html>

REGISTRATION

<http://www-cia.mty.itesm.mx/micai2000/registration.html>

TRAVEL INFORMATION

<http://www-cia.mty.itesm.mx/micai2000/travel.html>

TRAVEL & LOCAL ARRANGEMENTS

Leticia Rodriguez, ITESM-Mty, Mexico

Phone (52-8) 328-4377 Fax: (52-8) 328-4189

E-mail: lrodrig@campus.mty.itesm.mx

Conference Chairs
Francisco Cantu, ITESM-Mty, Mex
Enrique Sucar, ITESM-Mor, Mex



Program Chairs
English:
Osvaldo Cairo, *ITAM, Mex*
Luis A. Pineda, *UNAM, Mex*
Spanish:
Humberto Sossa, *IPN, Mex*
Jesús Favela, *CICESE, Mex*

Tutorial Chair
Marc Boumedine, ITESM-CCM, Mex

Workshop Chair
Humberto Sossa, IPN, Mex

lems by state-of-the-art satisfiability engines.

PROGOL—Developed by S. Muggleton, it is a system for inductive inference. It uses inverse entailment with mode declarations to derive the most specific clause with the mode language that entails a given example.

Summary

The general consensus of the attendees was that the field of LBAI has made significant contributions to both the theory and the practice of AI. In the theory of LBAI, contributions have been made to the foundations of commonsense reasoning, the formalization of non-monotonic reasoning, the development of semantics for theories, and results on computational complexity. In the practice of LBAI, we have seen the area of automated reasoning successfully applied to the verification of hardware systems and to improvements in solving satisfiability problems that have been incorporated into planning systems. The emergence of fast software systems implementing non-monotonic reasoning is the most significant development in this area in the recent years, highlighted by the systems SMOBELS, XSB, DLV, and DERES. System demonstrations of practical running systems that can be used for realistic knowledge base and non-monotonic reasoning applications, the development of planning systems, and tools for robotics were important contributions that were not available five years ago. Many long-term and short-term research efforts remain that must be handled before AI will be able to develop truly intelligent machines. The efforts in LBAI are expected to contribute to the development of intelligent machines. It is expected that a comprehensive report will be written concerning the workshop and that a book will be published that consists of papers from the workshop.

Jack Minker is a professor emeritus in the Department of Computer Science and the Institute for Advanced Computer Studies at the University of Maryland at College Park, Maryland. His e-mail address is minker@cs.umd.edu.

Case-Based Reasoning

Experiences, Lessons, & Future Directions

Edited by David B. Leake



Case-based reasoning is a flourishing paradigm for reasoning and learning in artificial intelligence, with major research efforts and burgeoning applications extending the frontiers of the field. This book provides an introduction for students as well as an up-to-date overview for experienced researchers and practitioners. It examines the field in a "case-based" way, through concrete examples of how key issues — including indexing and retrieval, case adaptation, evaluation, and application of CBR methods — are being addressed in the context of a range of tasks and domains.

ISBN 0-262-62110-X • 420 pp., index. \$42.00 softcover

AAAI Press • Distributed by The MIT Press • Cambridge, Massachusetts 02142

To order, call toll free:

(800) 356-0343 or (617) 625-8569.

MasterCard and VISA accepted.

www.aaai.org