



overall learning process. Kurland and Tenney (Chapter 5) closely examine the pivotal issues faced by designers of the HAWK intelligent tutoring system. Within this context they suggest numerous guidelines and meaningful examples from their work. Pirolli and Greeno (Chapter 6) provide a thorough analysis of the problem solving space for instructional design. The authors discuss critical design issues facing ITS developers, such as resources and constraints, in terms of levels-- from global (e.g., learning environments) and intermediate issues (e.g., topics, tasks), to local issues (e.g., texts).

Russell, Moran, and Jordan (Chapter 7) examine the instructional design environment from the perspective of IDE, a hypertext-based software environment. They describe the use of IDE to track instructional design decisions to create a rationale for a course design.

Macmillan, Emme, and Berkowitz (Chapter 8) focus on the important instructional planning process and suggest an architecture for dynamic instructional planning. The Self-Improving Instructional Planner (SIIP) is a generic blackboard-type architecture for the development of instructional planners within ITSs. The authors provide not only concise descriptions of SIIP and examples of its implementation, but also a valuable set of "lessons learned" that should be required reading by those integrating instructional planners into their ITS design.

The weak link in this section is Chapter 9, in which Wilkins, Clancey, Buchanan discuss differential modeling (i.e. isolating differences between observed behaviors by a problem solving agent and expected behavior based on a problem solving model). While it provides well-organized information and pointers to related work, it is more theoretical than previous chapters and may be less useful to those readers who learn best from other's experiences.

REPRESENTATION ISSUES: PART III Although the third section's theme is similar to that of previous texts and less interesting than the content of earlier sections of this text, the research presented here is useful in extending our current understanding. For example, Ritter and Feurzeig (Chapter 10) discuss the challenges involved in teaching "real-time tactical thinking," an area that stresses the already- problematic tasks of knowledge acquisition and student diagnosis. And although fault diagnosis has been a favorite area for ITS and maintenance system development Govindaraj (Chapter 11) describes a system that develops operators' metacognitive skills and knowledge in a manner that facilitates problems solving. Subsequent chapters within this section describe detailed development efforts and components of various types of systems (e.g., MACH-III).

SECTION IV: ARCHITECTURES FOR ITSs The final section is comprised of a series of chapters that describe theories and pragmatics involved in selecting or designing an ITS architecture. A common thread throughout these chapters is that although many lessons have been learned to date, many issues remain unsolved. Systems discussed include the Intelligent Maintenance Training System (Towne & Munro), QUEST problem solving (Feurzeig & Ritter), and BRIDGE (Bonar & Cunningham). However, Frye, Littman, and Soloway (Chapter 17) have created the most interesting chapter in this section, as they define the next wave of problems (e.g., designing the human interface, evaluation) in ITS development.

CONCLUSIONS Intelligent Tutoring Systems-Lessons Learned has something for every reader interested in intelligent tutoring systems. For those just entering the field, it provides the

"basics" in its coverage of knowledge acquisition, knowledge representation, and architectures. For those seeking to develop ITS that represent the best teaching practices, the text provides vital instructional design information. And for those of us hoping to move these promising systems into operational settings, it provides descriptions of "cutting edge research" and suggests areas for future focus.

New Books

Artificial Intelligence III

T. O'Shea

0-444-70508-2, 1988, 444 pages, \$94.75

Elsevier Science Publishing Co., Inc.

P.O. Box 1663

Grand Central Station, NY 10163

Mainstream AI research directions are reflected in the majority of the papers presented in this volume. Automatic reasoning and knowledge representation problems, inferencing, action planning and coordination, temporal reasoning, truth maintenance systems and machine learning are all represented. Other papers are concerned with logic and object-oriented programming, conceptual and implementation problems of AI programming, natural language processing, architectures and systems. The applications of AI are also considered, in the fields of software engineering, industry and the humanities, e.g. education, law.

Generating Natural Language Under Pragmatic Constraints

Eduard H. Hovy

0-8058-0248-7 (cloth), 1988, 232 pages, \$29.95

0-8058-0249-5 (paper), 1988, 232 pages, \$19.95

Lawrence Erlbaum Associates, Inc.

365 Broadway

Hillsdale, NJ 07642

What goes into making language? Building on the recognition that the generation of natural language is a goal-driven process, where many of the goals are pragmatic (i.e., interpersonal and situational) in nature, this work provides an overview of the role of pragmatics in language generation. Each chapter states some problem that arises in generation, develops a pragmatics-based solution, and then describes how the solution is implemented in PAULINE, a language generator that can produce numerous different versions of a single underlying message, depending on its setting.

Pragmatics and Natural Language Understanding

Georgia M. Green

0-89859-853-3 (cloth), 1988, 192 pages, \$24.95

0-8058-0361-0 (paper), 1988, 192 pages, \$12.50

Lawrence Erlbaum Associates, Inc.

This text provides both generalists in cognitive psychology and specialists in its various subdisciplines with a comprehensive and detailed introduction to the complex principles governing the interpretation of natural language. The author emphasizes the extent to which language interpretation and production depends on language users' assumptions and inferences about each other and

about particular aspects of the world. Compositional semantics, inferential pragmatics, and conventions of syntax are discussed.

Foundations of Intelligent Tutoring Systems

Martha C. Polson, J. Jeffrey Richardson

0-8058-0053-0 (cloth), 1988, 296 pages, \$39.95

0-8058-0054-9 (paper), 1988, 296 pages, \$19.95

Lawrence Erlbaum Associates, Inc.

A guide to the principles of ITS. Each chapter builds up the reader's knowledge of development, implementing, and evaluating an ITS. Basic research issues, state of the art systems and specific application projects are described by expert contributors.

The Nature of Expertise

Micheline T.H. Chi, Robert Glaser,

Marshall J. Farr

0-89859-711-0 (cloth), 1988, 464 pages, \$59.95

0-8058-0404-8 (paper), 1988, 464 pages, \$29.95

Lawrence Erlbaum Associates, Inc.

Over the last 20 years, expertise has become an important subject for investigation, due largely to developments in AI and cognitive psychology. *The Nature of Expertise* displays the variety of domains and human activities to which the study of expertise has been applied. This text reflects growing attention on learning and how expertise is acquired, by discussing the conditions that enhance and limit the development of high levels of cognitive skill.

Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference

Judea Pearl

0-934613-73-7, 1988, 552 pages, \$39.95

Morgan Kaufmann Publishers Inc.

P.O. Box 50490

Palo Alto, Ca 94303-9953

This book is an account of the theoretical foundations and computational methods that underlie plausible reasoning under uncertainty. The author provides a coherent explication of probability as a language for reasoning with partial belief and offers a unifying perspective on other AI approaches to uncertainty, such as the Dempster-Shafer formalism, truth maintenance systems, and nonmonotonic logic. The author distinguishes syntactic and semantic approaches to uncertainty - and offers techniques, based on belief networks, that provide a mechanism for making semantics-based systems operational. Specifically, network propagation techniques serve as a mechanism of combining the theoretical coherence of probability theory with modern demands of reasoning systems technology: modular declarative inputs, conceptually meaningful inferences, and parallel distributed computation. Application areas include diagnosis, forecasting, image interpretation, multi-sensor fusion, decision support systems, plan recognition, planning, speech recognition - in short almost every task requiring that conclusions be drawn from uncertain clues and incomplete information.

Practical Planning: Extending the Classical AI Planning Paradigm

David E. Wilkins

0-934613-94-X, 1988, 205 pages, \$34.95

Morgan Kaufmann Publishers Inc.

Planning, or reasoning about actions, is a fundamental element of intelligent behavior; and one that AI has found very difficult to implement. The most well-understood approach to building planning systems has been under refinement since the late 1960s and has now reached a level of maturity where there are good prospects for building working planners. *Practical Planning* is an in-depth examination of this classical planning paradigm through an intensive case study of SIPE, the most fully implemented planning system to date. The author, the developer of SIPE, defines the planning problem in general, explains why reasoning about actions is so complex, and describes all parts of the SIPE system and the algorithms needed to achieve efficiency. Details are discussed in the context of the problems and important issues in building a practical planner; discussions of how other systems address these issues are also included.

Representations of Commonsense

Knowledge

Ernest Davis

1-55860-033-7, Fall 1989, 500 pages, \$38.95

Morgan Kaufmann Publishers Inc.

A central problem in AI is how to endow computers with common sense. This requires systematizing a large body of knowledge about the commonsense world. This book suitable as a graduate text or reference, discusses AI techniques for representing and using knowledge in several, fundamental commonsense domains. Also discussed are algorithms, heuristics, and control structures that can be used to implement commonsense reasoning in these domains. The foundational issues of methodology, first-order logic and its extensions, and plausible reasoning are presented.

Reasoning With Incomplete Information

David Etherington

0-934613-60-5, 1988, 240 pages, \$22.95

Morgan Kaufmann Publishers Inc.

This monograph provides an extensive summary of work to date on nonmonotonic formalisms for reasoning without complete information about the world. Etherington imposes a unifying structure, based largely on semantic notions, on the widely disparate set of approaches found in the literature. Later sections of the book concentrate on Reiter's default logic and the various circumscriptive formalisms developed by McCarthy and others.

Readings in Knowledge-Based Systems

Randall Davis, Howard E. Shrobe,

and Peter Szolovits

1-55860-034-5, 1989, 500 pages (paper), \$29.95

Morgan Kaufmann Publishers Inc.

This Readings volume examines the important scientific and intellectual challenges in knowledge-based systems. Rather than focusing on specific systems and particular implementation problems, the editors have assembled this collection with a view

toward identifying the core scientific and engineering issues in knowledge-based system design. Each paper is introduced with a brief discussion of its contribution; the editors also provide an overall introduction to the volume. A challenging volume for designers, instructors, students, and professionals who use and advance the state of the art.

Automatic Refinement of Expert System Knowledge Bases

Allen Ginsberg

0-934613-96-6, 1988, 288 pages, \$22.95
Morgan Kaufmann Publishers Inc.

The research presented here contributes to the development of useful knowledge base refinement systems, both at the level of design, implementation and testing and at the level of methodology development for further research. Ginsberg generalizes and extends the empirically grounded heuristic approach to refinement generation developed by Politakis and Weiss and analyses strategies for using this approach in building automatic refinement systems. Further, he designs and implements a refinement metalanguage, RM. A provocative illustration of more powerful, robust, and self-improving expert system technology.

Automatic Program Debugging for Intelligent Tutoring Systems

William Murray

0-934613-98-2, 1988, 288 pages, \$24.95
Morgan Kaufmann Publishers Inc.

An exploration of the process by which student programs can be automatically debugged in order to increase their instructional capabilities. Murray provides a methodology and implementation for the diagnosis and correction of nontrivial recursive programs, with a case study of the automated debugger TALUS. Empirical evaluation of TALUS achieves high performance in debugging widely varying student solutions to challenging tasks.

Computational Learning Theory

David Haussler

0-55869-019-5, 1988, 400 pages, \$24.95
Morgan Kaufmann Publishers Inc.

A collection of original papers presented at the first Workshop on Computational Learning Theory, this volume brings together research from a variety of specializations focusing on rigorous and formal analyses of theoretical issues in machine learning. Many papers are revised and updated from the workshop presentations, incorporating results and insights from the workshop exchange.

Proceedings of the Fifth International Conference on Machine Learning

John Laird

0-934613-64-8, 1988, 467 pages, \$24.95
Morgan Kaufmann Publishers Inc.

Original articles presented by the recognized leaders in the field at the Fifth International Conference on Machine Learning, June 1988. Topics represented include empirical methods, explanation-based methods, genetic algorithms, connectionist learning, classification, probabilistic methods, problems solving,

SIGART Newsletter, January 1989, Number 107

language acquisition, machine discovery, and other related areas of research. These proceedings provide an up-to-date record of the state of machine learning research. This volume will be of interest to researchers and students in AI, cognitive science, philosophy of science, and education.

Designing AI-Based Software

A. Bahrami

0-20912-7, 1988, 205 pages, \$29.95
John Wiley & Sons, Inc.
P.O. Box 6792
Somerset, NJ 08875-9976

Provides readers with a practical selection of AI-based programming techniques that can be exported to every-day programming by exploring the capabilities of AI-based software and its uses in contemporary software engineering. Basic ideas about symbolic programming, problem and knowledge representation, search, control, problem solving, predicate calculus, and natural language Parser are presented. Also discussed is the latest frontier research in AI. Contains all of the important techniques for developing AI-based software in microcomputers. General audiences, as well as students of computer science, will find this book to be an invaluable guide to understanding current algorithmic and deterministic processes.

From Standard Logic to Logic Programming

A. Thayse

1-91838-5, 1988, 350 pages, \$37.95
John Wiley & Sons, Inc.

An introduction to the concepts and methods of AI, using a formal logical approach. Assuming no background in logic or AI, the book covers the fundamentals of mathematical logic, first-order theories, and classical logic, before describing non-classical logics. This is followed by discussion of a special type of knowledge representation - formal grammar - which can be interpreted as instructions of the logic programming language Prolog. Also describes Chomsky's hierarchy of formal grammars, the elements of which are shown to be connected with functional programming languages (such as LISP). The final chapter introduces Prolog and describes logic programming. Rules are stated for transforming logical representations into graphic and object representations.

AI A Handbook of Professionalism

Blay Whitby

0-21103-2, 1988, 273 pages, \$49.95
John Wiley & Sons, Inc.

An up-to-date overview of AI and its social consequences, answering such basic questions about AI as "what is AI?" and "who works in AI?" Addresses controversial issues that have been raised by the emergence and growth of AI, providing thoughtful discussion of the effects of funding from military sources, legal aspects of AI, the effect of AI in the workplace, and the Code for Professional Conduct for AI. The author stresses the importance and responsibility of AI developers in regards to the social impact of AI.

Coordination of Distributed Problem Solvers

Edmund H. Durfee

0-89838-284-X, 1988, 279 pages, \$55.95

Kluwer Academic Publishers

101 Philip Dr.

Norwell, MA 02061-1677

A valuable aid to researchers who want to understand the issues and the state-of-the-art in network coordination for distributed problem solving systems, **Coordination of Distributed Problem Solvers** introduces partial global planning as a new approach to coordination. Unlike previous distributed AI approaches, which are specialized for particular situations, this new approach provides a unified and versatile framework for dynamically coordinating independent nodes. This approach views control as a planning task. Nodes dynamically build and revise tentative plans about both local and network problem solving. Knowledge-based reasoning about coordination is seen as an integral part of a node's local decision making. A sample implementation for vehicle monitoring in a simulated distributed problem solving network is detailed. Further experiments show that partial global planning can be a flexible, practical approach for coordinating problem-solving networks and distributed computing systems.

Automating Knowledge Acquisition for Expert Systems

Sandra Marcus

0-89838-286-6, 1988, 288 pages, \$57.50 (cloth)

Kluwer Academic Publishers

Automating Knowledge Acquisition for Expert Systems describes a set of studies in automating knowledge acquisition that attempt to capture some of the expertise gained by knowledge engineers. The studies aim for creation of tools or methodologies designed to reduce the cost of building and maintaining expert systems. The tools discussed in this book have proved useful in creating and maintaining working expert systems. Because they are applicable to a range of tasks, the tools can be used to support multiple applications. They further the understanding of how task demands help shape the problem-solving strategies that can apply to them. Furthermore, the book's approach to creating these tools can be usefully employed by our programmer of expert systems to make the expert systems more accessible and maintainable.

Naive Semantics for Natural Language Understanding

Kathleen Dahlgren

0-89838-287-4, 1988, 272 pages, \$52.50 (cloth)

Kluwer Academic Publishers

Naive Semantics for Natural Language Understanding proposes a theory of word meaning representation, and describes its application in a functioning computational text understanding system. The theory, Naive Semantics, identifies word meanings with concepts, and sees concepts as rich, open-ended mental theories about the characteristics of objects and the implications of events. These native theories, which individuals hold, are generalizations about categories of objects and events, and are represented as typed features such as COLOR ("white") and

FUNCTION ("you wear it"). The features are not seen as primitives, but as other concepts, which could in principle number more than the number of English words.

Perspectives on Mind

Herbert R. Otto, James A. Tuedio

90-277-2640-X, 1988, 430 pages, \$79.00 (cloth)

Kluwer Academic Publishers

The analytical and continental traditions of Western philosophy are brought together in this work in a dialogue that explores the problems of mind in fresh and provocative ways. The original essays presented, accompanied by lively peer commentary, are organized thematically. The reflective essays emphasize a variety of perspectives on mind that have been developed by individuals having a variety of backgrounds and interests, but whose common central concern is an enhancement of our understanding of the nature of mind and intelligence. The work as a whole evokes new lines of thought and inquiry relative to the formulation and resolution of key issues in the philosophy of mind.

Logic Programming

Robert A. Kowalski and

Kenneth A. Bowen

Vol. 1: 61054-X

Vol. 2: 61055-8

Vol. 3: 61056-6

September 1988, 1700 pages, \$40.00 - each volume

\$75.00 two-volume set - softcover

The MIT Press

55 Hayward St.

Cambridge, MA 02142

These two volumes feature papers presented at the first joint meeting of the two principal logic programming conferences, held in August of 1988. The contributions cover all aspects of the field, including applications (particularly those that exploit the unique character of logic programming in AI, deductive databases, relations to other computational paradigms, language issues, methodology, implementations on sequential and parallel architectures, and theory.

Logic-Based Knowledge Representation

Peter Jackson, Han Reichgelt,

and Frank Van Harmelen

10038-X, November 1988, 300 pages, \$35.00

The MIT Press

This book explores the building of expert systems using logic for knowledge representation and meta-level inference for control. It presents research based on two hypotheses: that logic is a suitable knowledge-representation language, and that an explicit representation of the control regime of the theorem prover has many advantages. The chapters dealing with the representation problem present a reified approach to temporal logic that makes it possible to use nonstandard logics without extending the system, and describe a general proof method for arbitrary modal logics. Those dealing with the efficiency problem discuss the technique of partial evaluation and its limitations, as well as another possible solution known as assertion-time inference.

Kardio

Ivan Bratko, Igor Mozetic,
and Nada Lavrac

02273-7, November 1988, 200 pages, \$35.00

The MIT Press

A successful complex expert system, Kardio's model of the human heart is designed for the diagnosis of cardiac arrhythmias. Kardio is highly skilled in the reading of ECG recordings, and can be used as a diagnostic tool in ECG interpretation. In this groundbreaking account, the authors show how the model was compiled, by means of qualitative simulation and machine learning tools, into various representations that are suited for particular expert tasks. They investigate a hierarchical organization of a qualitative model and outline an experiment whereby the construction of a deep model is automated by machine learning techniques. The book contains a complete model of the electrical system of the heart.

Ontic

David A. McAllester

13235-4, September 1988, 375 pages, \$37.50

The MIT Press

ONTIC, the interactive system for verifying "natural" mathematical arguments that David McAllester describes in this book, represents a significant change of direction in the field of mechanical deduction, a key area in computer science and AI. ONTIC is an interactive theorem-prover based on novel forward-chaining inference techniques. It is an important advance over earlier systems for checking mathematical arguments as Automath, Nuprl, and the Boyer-Moore system. The first half of the book provides a high-level description of the ONTIC system and compares it with other automated theorem-proving and verification systems. The second half presents a complete formal specification of the inference mechanisms used.

McAllester's is the only semi-automated verification system based on classical Zermelo-Fraenkel set theory. It uses object-oriented inference, a unique automated inference mechanism for a syntactic variant of first-order predicate calculus. The book shows how the ONTIC system can be used to check such serious proofs as the proof of the Stone representation theorem without expanding them to excessive detail.

Articles

Abstractions in Semantic Networks: Axiom Schemata for Generalization, Aggregation and Grouping

Ulrich Schiel

Departamento de Sistemas e Computacao
Federal University of Paraiba

Av. Aprigio Veloso s/n
58100 Campina Grande
BRAZIL

One of the most common structures used to represent

knowledge are the semantic networks. They were also used in the database field for the semantic data modeling and development of conceptual schemata. A semantic network is a graph whose nodes represent concepts, classes or sets of entities and the arcs model relationships between these concepts. Besides the named relationships obtained from the modelling of an application with a semantic network, some special, structural relationships between concepts can occur and have a rich expressive power. They are known in the database field as abstraction hierarchies and called *Generalization, Aggregation and Grouping* [SS-77, HM-78] and correspond to the semantic links IS-A, PART-OF and ELEMENT-OF respectively. In order to allow a correct use of these links and their interrelationship, we present here a set of axiom schemata formalizing these concepts.

A recent related work is [Ga-88]. But he does not distinguish between PART-OF and ELEMENT-OF nor distinguish between IS-A and INSTANCE-OF and does not give axioms relating the semantic links together.

DEFINITIONS We distinguish between *concepts* or classes and their *instances* or elements. In order to refer to instances in most cases we need an identification of the instance. We use the notation (x) to mean an instance identified by x. Therefore "John" is the string John but ("John") is an entity identified by the string John. In order to determine of what kind of concept ("John") is we write ("John"):PERSON or ("John"):STRING:PERSON. We use small letters for variables of instances or entities and capital letters for concept-variables and examples of concepts.

The three structural links between concepts can be defined in several frameworks (set theoretic, logic, etc., see [Br-83] for the IS-A relationship). We give a set theoretic definition of the abstractions [Sc-83] and introduce a set of predicates in order to specify these structural relations between concepts. If C is a concept then |C| means the set of instances of C.

1. A concept G is a *generalization* of a concept C iff $|C| \subseteq |G|$.
2. A concept A is an *aggregation* of concepts $C_1, \dots, C_n$ iff $|A| \subseteq |C_1|$.
3. A concept G is a *grouping* of a concept E iff $|G| \subseteq P(|E|)$

PREDICATES The first predicate expresses the relationship between instances and concepts:

instance-of(e,C) : e is an instance of the concept C.

E.g. instance-of("John",STRING)

instance-of("John"),PERSON)

The instance-of predicate establishes the link between instances and concepts. All the other, which are the logical counterpart of the set-theoretical definitions below, can be classified into two categories:

a) Predicates between concepts

is-a(S,G): G is a generalization of S

E.g.: is-a(SAILBOAT, BOAT)

is-a(TRIANGLE, POLYGON)

In some cases it is of interest to consider a role that determines the semantics of the generalization/specialization. In this case we write is-a(S,G,r). E.g. is-a(MALE,PERSON,sex). With several roles we can define distinct partitions of a class. E.g.

sex(PERSON)=MALE,FEMALE

age(PERSON)=YOUNG,MIDDLE,OLD.

part-of(C,i,A): A is an aggregate class and C is the i-th component of A. If the position of the component is not of interest,