T. Christaller
Konsistenzpruefungen bei sich veraendernden Wissensbasen
(Consistency-checking in Changing Knowledge Bases)

W. Emde, C. Habel, C.R. Rollinger
Automatische Akquisition von inferentiellem Wissen
Automatic Acquisition of Inferential Knowledge)

C. Habel
Zur Repraesentation indifinit deskribierter Objekte in KIL-Systemen
(On the Representation of Indefinitely Described Objects in AI Systems)

K.v. Luck, B. Owsnicki
N.N: A View of Planning in Chess

UNIFIKATION UND DEDUKTION
(Unification and Deduction)
J. Siekmann, P. Szabo
Universal Unification
(Invited Lecture)

W. Benda, G. Hornung, P. Raulefs, F. Vollmann
Der META-Beweiser fuer die Zahlentheorie
(The META-prover for Number Theory)

M. Bidoit
Proofs by Induction in *Fairly* Specified Equational Theories

M. Halpern
Inductive Inference in Finite Algebraic Structures

A. Herold
Universal Unification and a Class of Equational Theories

G. Smolka
Completeness of the Connection Graph Proof Procedure for Unit-Refutable Clause Sets

NATUERLICHSPRACHLICHE SYSTEME
(Natural Language Systems)

M. Gehrke
Rekonstruktion aufgabenorientierter Dialoge mit einem mehrstufigen Parsing-Algorithmus auf der Grundlage kaskadierter ATNs
(Reconstruction of Task-oriented Dialogues with a Multiple-level Parsing Algorithm Based on Cascaded ATNs)

W. Hoeppner
ATN-Steuerung durch Kasusrahmen
(Case-driven ATN Parsing)

M. Hussmann
Ellipsenrekonstruktion im Kotext: Forderungen an eine Ellipsenkomponente und ihre Realisierung in SWYSS
(Ellipsis Reconstruction in Co-text: Requirements for an Ellipsis Component and their Actualization in SWYSS)

C.B. Schwind
Natural Language Access to PROLOG Database Systems

BOOKREVIEW

Artificial Intelligence: An MIT Perspective, Volume 1: Expert Problem Solving, Natural Language Understanding, Intelligent Computer Coaches, Representation and Learning

Reviewed by:
Varol Akman
Electrical, Computer and Systems Engin. Dept
Rensselaer Polytechnic Institute
Troy, New York 12181

(Edited by Patrick Henry Winston and Richard Henry Brown, The MIT Press, Cambridge, Massachusetts, 1979, 492 pages, $12.50 paperback, ISBN 0-262-73058-8.)

Artificial Intelligence: An MIT Perspective consists of two volumes. Volume 1 covers the areas cited in the title of this review. Volume 2 (which will be the subject matter of an upcoming review) has sections on Understanding Vision, Manipulation and Productivity Technology, Computer Design, and Symbol Manipulation.

In these volumes, Winston and Brown aim to review almost all branches of AI research conducted at MIT and to provide a good starting point for anyone (novice or expert) who really wants to obtain a clear perspective of MIT AI Lab's outstanding contributions. With one exception, all of the papers in Volume 1 are MIT AI Lab memos. Most of them also have been published in AI-related journals and symposium proceedings.

Due to space limitations, the editors state in the Introduction that these books do not attempt to be exhaustive. Selected articles are only representative of the surveyed research areas. To encourage cover-to-cover reading, the sections are often kept technically simple and descriptive. Thus, a sophisticated reader would not be able to find recipes for "how to implement things" but rather neat accounts of "what can be achieved." In the editors' words, the reader is invited to treat these books "as a collection of hors d'oeuvre" which may be complemented by "entrees accessible through the references" at the end of each paper.

Surprisingly (and luckily) enough these books have the appearance of a good textbook in spite of the fact that they are basically a collection of papers. Each paper is preceded by a synopsis which also serves as a historical overview and explains the importance of the reported work. This is very encouraging for the intermediate reader since initial unfamiliarity with the topic under consideration is skillfully lessened and passages between consecutive papers by different authors become smoother. Even in this respect, Artificial Intelligence: An MIT Perspective stands as a "a job well-done" which must be studied thoroughly by anyone who thinks of editing a similar book on the subject.

Although these volumes were prepared directly from camera-ready computer-generated material, the quality of printing is excellent for joyful reading. The illustrations are perfect and there are only a few (unimportant) typos. Each volume is followed by a wide subject index. Unfortunately, the editors apparently did not think that an author index would also be useful.

In the remainder of this review, I will give a brief description on each paper in Volume 1.

Section I: Expert Problem Solving consists of five articles. They generally introduce problem solving programs which employ a compact (and interacting) database of facts to handle problems in a specialized domain.

- J. de Kleer opens this section with his paper, "Qualitative and Quantitative Reasoning in Classical Mechanics." This is a discussion of a program (NEWTON) which solves simple mechanics problems. This paper appeared in Proceedings of the Fifth International Joint Conference on Artificial Intelligence (1977).
- "Problem Solving About Electrical Circuits," by R.M. Stallman and G.J. Sussman is a step toward understanding complex electronics circuits. The authors describe a mechanism which makes assumptions and gains perceptiveness from contradictions arising from assumptions. This paper was published in Artificial Intelligence (1977).
- The third paper o this section, "Explicit Control of Reasoning," is by J. de Kleer, J. Doyle, G.L. Steele, and G.J. Sussman. It presents a means for embedding expert knowledge about expert knowledge in a database. It also provides a technique for describing reasons for belief in an assertion. Accurate dependencies allow correct assignment of responsibility for invalid beliefs.
- "A Glimpse of Truth Maintenance," by J. Doyle is about maintaining a database system of currently valid deductions. Such a self-contained subsystem relieves a higher-level (parent) expert of this responsibility. This paper is a condensed version of Doyle's MIT MS thesis.
- The last paper of this section, "Design of a Programmer's Apprentice," is by C. Rich and H.E. Shrobe and reports on the theoretical basis to support the design of a programmer's apprentice (a system which cooperates with a programmer during the life-cycle of a program).

Section II: Natural Language Understanding and Intelligent Computer Coaches consists of five articles. These papers describe research done to help computers understand everyday English and exhibit a human-like understanding of useful support tasks.

- M.P. Marcus opens this section with, "A Theory of Syntactic Recognition for Natural Language." In this overview of his MIT Ph.D. thesis (published by the MIT press under the same title), Marcus demonstrates an alternative to augmented transition networks in modeling syntax in natural language understanding.
- C. Sidner's, "Disambiguating References and Interpreting Sentence Purpose in Discourse," offers a frame-based theory to handle omissions, referential terms, etc. in everyday usage of English. This paper appeared in Proceedings of the Second National Conference of the Canadian Society of Computational Studies of Intelligence (1978).
- "Using Frames in Scheduling," by I.P. Goldstein and B. Roberts provides another frame-based approach (and a knowledge-based program called NUDGE) to resource allocation tasks such as meeting scheduling and high-level planning. This paper is in Proceedings of the Fifth International Joint Conference on Artificial Intelligence (1977).
- The fourth paper of this section, "Developing Support Systems for Information Analysis," by J.L. Stansfield, is an inquiry into the problems to be solved to make a computer knowledgeable enough to carry out common-sense news analysis. Stanfield's program is specialized to the international wheat market domain.
- In the last paper of this section, M.L. Miller and I.P. Goldstein study, "Planning and Debugging in Elementary Programming." This paper deals with explanation and user modeling issues in the domain of coaching novice programmers. It appeared in International Journal of Man-Machine Studies (1981).

Section III: Representation and Learning is a survey of problems encountered in describing things with a set of conventions. A good representation is the foundation for all areas of AI. There are four articles in this section.

- In, "Learning by Creating and Justifying Transfer Frames," P.H. Winston argues that learning theory is strongly dependent on M. Minsky's theory of representation via frames. This paper appeared in Artificial Intelligence (1978)
- The second paper is, "Description and Specialization of Concepts," by W.A. Martin. In this paper, Martin formulates his theory of representation which is specially tailored to deal with language understanding issues. This is the only paper in this collection which was published as a MIT LAB for Computer Science memo.
- M. Minsky's, "The Society Theory of Thinking," establishes a theory of presentation which is more on the procedural side. In this representation, societies of simple agents intercommunicate to fulfill intelligent tasks. This paper was published in Proceedings of the Fifth International Joint Conference on Artificial Intelligence (1977).
- In the last paper of this section (and Volume 1), "Representing and Using Real-World Knowledge," S.E. Fahlman studies the issues related to networks of real-world concepts. Using a parallel network of simple processors, he concisely represents and processes knowledge. This is an overview of his MIT Ph.D. thesis published by the MIT Press with the title, "NETL: A System for Representing and Using Real-World Knowledge."

In summary, this two volume set is an excellent overview of AI as theorized and practiced at MIT and one may easily get acquainted with "more AI" by simply following the references at the end of each paper. Although all the articles in Volume 1 (and similarly in Volume 2) are individually available in AI journals and symposium proceedings, their structured and coherent presentation under general headings make these two volumes indispensable. I strongly believe that every serious AI researcher should have this set available for reference and for reading. I will be very pleased if Winston and Brown possibly spare some time to update and enlarge this beautiful work with new contributions once in a while. (Surely, things have changed quite a bit since 1979). I can even hope that his work of Winston and Brown will be the origin of a much wider and larger set of books embracing all memos published by the MIT Lab since its inception. Then, one can call that work the, "MIT Encyclopedia of Artificial Intelligence."

THE HEDONISTIC NEURON: A THEORY OF MEMORY, LEARNING, AND INTELLIGENCE

Stephen W. Smoliar
Schlumberger-Doll Research
P. O. Box 307
Ridgefield, CT 06877

(A. Harry Klopf, Hemisphere Publishing Corporation, Washington, 1982, 140+xvii pp.)

> Science is a wonderful thing, but it has not succeeded in maximizing pleasure and minimizing pain, and that's all we asked of it.
>
> *Notes and Comment,*
> **The New Yorker,**
> June 13, 1970
> cited by Paul Dickson,
> in **The Official Rules**

In an age in which verbose obfuscation becomes more fashionable with every new government appointment, one has to appreciate an author who can give a cogent summary of his entire book in the second paragraph of his Preface:

> The book examines a theory of brain function based on the postulate that neurons "seek" excitation and "avoid" inhibition, this goal being pursued within well-defined limits that preclude an epileptic outcome (except under pathological circumstances that, indeed, we know do arise). In the development of the theory, it will be seen that a plausible adaptive mechanism exists that can account, in mechanistic terms, for the postulated neuronal behavior. Also, the notion of an excitation-seeking, inhibition-avoiding neuron will be shown to be consistent with experimentally observed neuronal behavior. Neuronal and cortical polarization studies, the mirror focus, and epileptic foci appear to be understandable in light of the proposed neuronal model. At a psychological level, habituation, dishabituation, classical and operant conditioning, and extinction can be shown to be straightforward consequences of a goal-seeking neuron like the kind proposed. None of this is to suggest that the validity of the proposed theory is demonstrated here. Far from it. Difficult experiments will have to be performed at the neuronal level to test the theory rigorously. Until these experiments are accomplished, the significance of the theory lies in its offering a fundamentally new view of brain function, a view that suggests alternative, and perhaps more productive, experiments.

However outrageous one finds the thesis (and if nothing else, the thesis is certainly unorthodox), one must give the author credit for his honesty. If he seriously intends to pursue the experimental program proposed at the end of the above paragraph, then one must be fair enough to wait for the results to come in. In the meantime, this little book provides an interesting source of provocation for those whose views of cognition may have fallen into a rut.

One might, perhaps, question whether hedonism is an appropriate metaphor for brain research. Viewers of the NATURE series on animal behavior may tend to be a bit cautious about using any human characteristic as a general metaphor for behavior, and they certainly have been given evidence that unadulterated behaviorism is not necessarily "the way" of animal behavior. Nevertheless, one scientist should never attempt to tell another scientist what his hypotheses should be, particularly if that latter scientist has a well-formed plan of experimental investigation. One can only wonder, however, if the author will be able to provide adequate support for his research plan unless he is doing research in the best of all possible worlds.

Given that I am not very well versed in neurophysiology, I am not in a position to analyse the author's arguments in great detail. As a computer scientist, however, I must admit to a certain degree of disappointment. The initial discussion is concerned with the development of a neurophysiological model, and this model is built in such a way as to accommodate the myriad features discussed in the quoted paragraph. Unfortunately, the model is so big and so general that a reader who is armed with little more than commen sense cannot really develop a feel for what it actually does. It would be nice to see some sort of concrete demonstration of the model, something which would justify the author's faith in it. After all, the one danger of the book is that the reader will forget the author's own disclaimer. The line between a valid proposal and an idle speculation is often hard to draw.

One might also fault the author for setting up outdated research in artificial intelligence research as a straw man. For example, he feels that a major fault of AI is that it has been concerned with systems built on von Neumann-style architectures. This may have been true when he first wrote that sentence, but it is certainly no longer true. He also seems to feel that the concept of "knowledge base" involves "a passive approach to learning and memory," an assertion which is also questionable as of the writing of this review.

In summary, one should not expect too much from this book. It neither promises nor delivers a panacea. It does offer a lot of provocation. Given that it entails less than 200 pages of reading matter, one could do far worse than devote one's attention to it.

Artificial Intelligence: An MIT Perspective, vol. 2

Reviewed by:
Ram Nevatia
Dept. of Electrical Engineering
University of Southern California
Los Angeles, CA 90089-0272

(Edited by P. H. Winston and R. H. Brown MIT Press 1982 (paperback edition))

This is a paperback edition of the original version published in 1979. Vol. 2 is divided in three sections dealing with vision, manipulation and machine design. (Vol. 1 deals with expert problem solving, natural language, and representation and learning.) Each section contains chapters written by various MIT researchers. In general, each chapter describes recent (as of the original date of publication) MIT research in the particular areas. Each section also provides a nice overview and summary by the editors.