Second Edition

ENCYCLOPEDIA OF ARTIFICIAL INTELLIGENCE

Volume 1
A–L

Awarded
American Library Association's
Outstanding Reference Source
Association of American Publishers Award
Best New Professional and Scholarly Publication

Stuart C. Shapiro
Editor-in-Chief

Cognitive Science

Zenon W. Pylyshyn, COGNITIVE SCIENCE, 188-192

From Shapiro, Stuart C., Editor-in-Chief, Encyclopedia of Artificial Intelligence, 2nd Ed., John Wiley & Sons, Inc., NY, 1992.

COGNITIVE SCIENCE

RELATION TO OTHER FIELDS

Cognitive science is an emerging field of study, boundaries of which are far from being well defined. A report prepared for the Alfred P. Sloan Foundation (a portion of which is reproduced as an appendix to Pylyshyn, 1983) defines it as "the study of the principles by which intelligent entities interact with their environments" and notes that "by its very nature this study transcends disciplinary boundaries." In particular, the distinctions among cognitive psychology (qv), AI, and cognitive science are extremely blurred in practice. This blurring is additionally exacerbated by the fact that research that clearly qualifies as cognitive science is being done in academic departments (as well as government and industrial research laboratories), the titles of which identify them with disciplines as diverse as psychology, computer science, linguistics, anthropology, philosophy, education, mathematics, engineering, physiology, and neuroscience. From an informal survey of cognitive science publications, it is shown that papers in cognitive science journals cited other papers in a very wide range of fields (Pylyshyn, 1983).

Cognitive science is also extremely closely related to AI. When the editors of the journal *Artificial Intelligence* decided to help their readership keep up with some of the literature in closely related disciplines by publishing regular "Correspondent's Reports" on work in these fields, they selected the areas of philosophy and logic, robotics, software engineering, natural language, cognitive psychology, and vision. Of these, all but perhaps software engineering and parts of robotics would be considered core areas of cognitive science research. Indeed, it has been argued (Newell, 1970, 1973; Pylyshyn, 1979, 1980) that AI and cognitive science may be nothing more than two paths to the same end: understanding the nature of intelligent action in whatever physical form it may occur. The difference between them, according to this view, consists mainly in research style: AI takes the high road of asking how instances of intelligence can be realized (ie, how they are possible) within the constraints of known computational mechanisms or how they might be attainable by the design of new mechanisms (ie, new computational architectures), whereas cognitive science places greater emphasis on the question of how instances of intelligence are in fact realized within one particular architecture, the one constituted by the human mind. Because of this difference in orientation, many experimentally oriented cognitive scientists tend to place a somewhat greater premium on empirical fit, on testing processes against psychological data to determine not only whether the two are input–output equivalent but also whether they are strongly equivalent, that is, whether in both cases the behavior is produced by the same information-processing means. The notion of strong equivalence is central to much cognitive science, although it is not often discussed explicitly. According to one interpretation (Pylyshyn, 1984), two processes can be strongly equivalent only if they produce the same behavior using the same computational process (or algorithm) and the same symbolic representations, something that is possible only if the two systems have functionally identical computational architectures (ie, the same primitive operations, the same resource constraints, and the same symbolic notation).

Despite this difference in principle between cognitive science and AI, differences in practice are minimal. Indeed, it has even been argued (Pylyshyn, 1979) that a convergence of the two approaches may be inevitable inasmuch as both adhere to a notion of intelligence that is inherently anthropocentric or human relative, at least at the present time.

Of course, the two fields diverge considerably in their applied side. A great deal (though by no means all) of applied cognitive science deals with such problems as designing better human–machine interfaces (see HUMAN–COMPUTER INTERACTION), better pedagogical methods (see EDUCATION, AI IN) better communications techniques, better aids for the disabled (see PROSTHESES), or better methodologies for discovering such useful things as what experts know (see KNOWLEDGE ACQUISITION) or why children fail to read or do mathematics. What identifies these as cognitive science rather than simply applied psychology investigations is the fact that they take a fundamentally computational view of the nature of the cognitive process involved; they view cognitive process as consisting of the execution of symbol manipulation procedures. Although it is clear that the fruits of such pursuits are relevant to what people in AI do, the work itself frequently requires different skills and proceeds using different methodologies than are typically (though, again, not always) found in AI laboratories.

In contrast to this approach, applied AI places heavy emphasis on finding a practical match between available computational techniques and applications crying out for solution. As in all engineering or applied technology pursuits it must find suboptimal solutions to practical problems and proceed by incremental refinement. In terms of what has been referred to as the power generality tradeoff (Ernst and Newell, 1969), applied AI must perforce settle for the power end of the dimension. But none of this need be true, and indeed generally is not true, of basic research in either AI or cognitive science, where the overlap is great enough that many are tempted to view AI as the more theoretical and more formal end of the spectrum of cognitive science research.

This still leaves the question: what is cognitive science? If it is simply the attempt to understand mental activity (or, as in the earlier quote from the Sloan report, to understand how intelligent entities interact with their environments), how is it different from psychology, especially from that branch of psychology that studies thinking, perception, memory, language, and so on, that is, cognitive psychology (qv)? Many people believe that cognitive science represents a new paradigm for understanding cognition, a paradigm that clearly owes much to developments in computer science. Yet a better characterization is desired, because if it is a new paradigm, it would be useful to know how it differs from other paradigms and on what assumptions it stands. It would be helpful to know this both in the abstract (ie, what are some distinguishing principles of cognitive science?) and in terms of concrete examples of how the new science is practiced and what it is seen as accomplishing, or at least trying to accomplish.

Many attempts at a statement of what cognitive science is have been made. One of the earliest was the unpublished report prepared by a committee (under the editorship of George Miller) for the Alfred P. Sloan Foundation from which the earlier quote was taken. This report characterizes what is special about cognitive science and what runs through all the diverse work that falls under its scope by defining its research objective as being "to discover the representational and computational capacities of the mind and their structural and functional representation in the brain." Although extremely general, this represents a fair statement. In addition to this early statement, the journal *Cognitive Science*, the official organ of the Cognitive Science Society, has published a number of articles that attempt to characterize the field, beginning with its initial editorial, and has included a number of papers first presented at the inaugural conference of the Cognitive Science Society in 1979 (these were published in several issues of the journal, beginning with volume 4, 1980). An attempt at a systematic argument that cognitive science is not just a marriage of convenience but a genuine field of study has been published (Pylyshyn, 1984).

This article provides examples of the kinds of problems that cognitive scientists are interested in pursuing and the approaches that they take, pointers to literature that gives further details of such examples, and a brief statement of why some people believe that cognitive science is not just a collection of research problems that in one way or another are concerned with reasoning but a genuine scientific domain or inquiry. It should be noted, however, that this review is not without personal bias. It is primarily an attempt to characterize cognitive science rather than to catalog its current research directions, which are likely to change radically in the next few years in any case. Moreover, this article presents a view of what is constitutive of cognitive science which the author believes to be correct and borne out by the classic work in the field (and which is defended at some length Pylyshyn, 1984), yet this view nonetheless flies in the face of claims being made by some people who are legitimate researchers in cognitive science. This view concerns the symbol-processing nature of cognition (what is referred to below as the representational metapostulate). Although this is not the proper forum for a debate on such issues, it is believed that the notion of symbolic representation is so very central to cognitive science and continues to be the central theoretical assumption underlying virtually all work in the field that it is appropriate to lay it out explicitly, even in the sketchy form presented here.

SOME EXAMPLES OF COGNITIVE SCIENCE RESEARCH PROBLEMS

Language

The study of the human capacity for language is one of the oldest areas of research in cognitive science. It is also one that has changed dramatically in the past two decades, partly under the influence of formal linguistics and partly because of attempts to develop computer systems for understanding natural language (see NATURAL LANGUAGE UNDERSTANDING). It thus provides a prime example of cross-disciplinary cognitive science research, albeit one that continues to be steeped in controversy. In recent years this study has also encompassed work by philosophers, as researchers become more concerned with issues of semantics and pragmatics, with problems of meaning and discourse that had occupied philosophers long before these problems arose in AI. It also brought in the work of clinical neuroscience, which investigated the taxonomy of language deficits caused by trauma and disease.

This work has led to computational models of language performance. At the present time a number of alternative models of syntactic analysis (see PARSING) have been published and psycholinguistic research provides provocative evidence that parsing proceeds with only minimal input from the rest of the cognitive system, as is also the case, by the way, in most computational language understanding systems; a notable exception is the work of Schank and his colleagues (Schank and Abelson, 1977). Experimental studies have also shown clearly that the lexical lookup (see MORPHOLOGY) phase of grammatical analysis retrieves many homographs or homonyms of ambiguous words (Swinney, 1979; Seidenberg and Tanenhaus, 1985), thus empirically validating one computational proposal.

Vision

The idea, popular in the 1950s, that perception consists of hypothesis testing was challenged first by people working on computational vision (see EARLY VISION) (Zucker and coworkers, 1975; Marr, 1982), who argued that it would be highly wasteful to not extract as much information as possible from the initial image before bringing cognitive processes to bear. Some models (Marr, 1982) showed that a considerable amount of processing could be done in a data-driven manner (see PROCESSING, BOTTOM UP AND TOP DOWN). These ideas were then validated by psychophysical investigations as well as by findings from neuroscience [eg, concerning the existence of separate spatial frequency channels, motion detectors (see VISUAL MOTION ANALYSIS), sensitivity to maxima in intensity derivatives, etc]. Some of this cross-fertilization has been nicely illustrated (Brady, 1981). Although this work is described in some detail elsewhere in this encyclopedia, it is in fact an excellent example of cognitive science research that falls at the more computational end of the spectrum. The relevance of both the vision and the psycholinguistics work to the understanding of mind has been discussed in an insightful and provocative way (Fodor, 1983).

Expertise and Qualitative Reasoning

The study of expert systems (qv) (or, as it is sometimes called, knowledge engineering) both inspires and benefits from experimental investigations of how experts in such areas as physics, mathematics, electronics, medicine, or chess differ from their inexperienced counterparts. Findings concerning how experts structure their knowledge and how this structure differs from that of less experienced performers is an interesting chapter in recent cognitive science. These investigations also relate to studies in both psychology and AI of how people reason by building qualitative mental models (qv) (Hobbs and Moore, 1984; Gentner and Stevens, 1982; Johnson-Laird, 1983).

Models of Human Performance in Various Tasks

In this category is found computational models of human performance on arithmetic (Brown and Van Lehn, 1980), tasks involving interacting with text editors (Card and coworkers, 1983), typing and other skills (Cooper, 1983), and reasoning with spatial problems (Kosslyn, 1980). Closely related to this work is the general study of cognitive skill, its acquisition, and its nature (Anderson, 1982). Understanding cognitive skill requires distinguishing cognitive capacities from performance differences that arise from differences in knowledge or habit, a difference that parallels the distinction between functional architecture and computational procedures. The importance of this distinction to understanding the nature of cognitive processes (and of strong equivalence) has been discussed (Pylyshyn, 1984, 1981).

Learning

The area of learning was one of the most thoroughly investigated during the last half century of psychology, with very little progress on what people call learning in everyday life. The work was guided by preconceived ideas about the underlying mechanism (namely, association) rather than by a careful analysis of the types of learning and the types of mechanism capable of meeting the sufficiency condition that is central to cognitive science. More recent work on language learning by cognitive scientists has shown that the acquisition of syntax from the kind of evidence generally available to the child would not be possible without severe constraints on both the structure of the languages that can be learned and severe constraints on the mechanisms that could learn such languages. In particular, it is necessary that the range of grammars that the organism could consider as possible hypotheses must be extremely limited (Wexler and Cullicover, 1980). The same may also be true of concept acquisition (Demopoulos and Marras, 1985). More recent work on learning within AI has also provided new ways to look at some forms of learning in humans (Michalski and co-workers, 1986). (see also LEARNING, MACHINE).

CONCLUSION: SOME CHARACTERISTICS OF COGNITIVE SCIENCE

Cognitive science is not the only form in which the search for an understanding of mind is proceeding. What characterizes this particular class of approaches is an allegiance to the network of ideas that might roughly be summarized as follows (Pylyshyn, 1983).

1. The approach is formalist in spirit: that is, it attempts to formulate its theories in terms of symbolic mechanisms of the sort that have grown out of symbolic logic (qv), although the apparatus of formal logic itself very rarely appears in cognitive science theories.

2. The level of analysis, or the level at which the explanations or theories are cast, is functional and is described in terms of its information flow. What this means in particular is that this approach factors out questions such as how biological material carries out the function and how biochemical and biophysical laws operate to produce the required information-processing function. This factorization is analogous to the separation of electrical engineering considerations from programming considerations in computer science. This does not mean that questions of biological realization are treated as any less important, only that they represent a distinct and to a large extent independent area of study. According to this view, neuroscience contributes an understanding of how such computational processes as are uncovered by empirical observations of human capacities are realized by biological mechanisms.

Not everyone agrees that cognition can be studied independently of its neurophysiological instantiation. There is, for example, an approach, sometimes called connectionist (see CONNECTIONISM), which attempts to build models of cognition that are guided more closely by ideas from neuroscience than by symbol-processing ideas from current computer science. Some examples of such models have been published (Anderson and Hinton, 1981; *Cognitive Science* **9**(1), 1985). Although this approach is extremely promising from the perspective of modeling the functional architecture of the mind, there is considerable doubt that it can displace rule-governed symbolic processes entirely, as some have claimed (Pylyshyn, 1984).

3. In addition to factoring apart questions of capacities from questions of biological realization, the approach is also characterized by the techniques it uses in formulating its theories and in exploring the entailments of its assumptions. The most widely used (although not universal) technique is that of computer implementation. Thus an important methodological goal of cognitive science is to specify symbol-processing mechanisms that can actually exhibit aspects of the behavior being modeled. Adherence to such a sufficiency criterion makes this approach in many respects like a design discipline rather than natural science, at least insofar as the latter typically attempts to uncover a small set of fundamental axioms or laws. Its concern with synthesis makes it one of the "sciences of the artificial," (Simon, 1969) along with AI.

4. The approach tends to emphasize a strategy sometimes referred to as top-down analysis, in which a premium is given to the task of understanding how the general cognitive skill in question is possible (consonant with the constraint of mechanism) in contrast with the task of accounting for empirical particulars. This difference in style contrasts with the traditional approach in experimental psychology that emphasizes the observational fit of models. The contrast has been examined (Pylyshyn, 1979; Newell, 1970; Sloman, 1968).

5. This commitment to the informational level also places the enterprise in contrast to the phenomenological approach in which the existential notions of significance, meaningfulness, and experiential content are given a central role in the analysis and with behaviorism, which attempts to analyze behavior without appeal to internal representational states. These issues have been discussed (Pylyshyn, 1984; Cummins, 1983; Fodor, 1981; Haugeland, 1981; Dennett, 1979).

The above general characteristics of cognitive science are also shared to various degrees by other scientific disciplines. The formalist or symbolic mechanistic character 1 is deeply entrenched in contemporary linguistics (especially in generative grammar), decision theory, and even parts of anthropology (eg, Levi-Strauss). The functionalist perspective 2 is now quite general in psychology and philosophy of mind as well as in engineering, where it is referred to as the black-box approach. Both 1 and 2 are fundamental to computer science as well as to any science that concerns itself with notions such as the flow of information or the distribution of control. Such ideas have thus affected everything from engineering to management science and even political science (Deutsch, 1963). Criteria 3 and 4 are not quite so prevalent as the first two. For example, the desire to synthesize aspects of the phenomena being modeled as part of the attempt to understand it is not widespread in the social sciences outside of the areas of cognitive psychology and management science [especially the branch of the latter called industrial dynamics (Forrester, 1971)], nor is it yet very common in biology [see, however, Marr's critique (1975) of theories in neurophysiology that fail to characterize the constructive computational aspect of biological function]. Even modern linguistics, which is in many ways a prototypical cognitive science, places little emphasis on the human capacity to actually generate samples of performance. However, examples of the contrary trend are available (Marcus, 1979; Berwick and Weinberg, 1984).

The Representational Metapostulate

Although, as suggested earlier, there are a number of theoretical and methodological characteristics that pervade a variety of approaches to the understanding of intelligence and human cognition, there is one overriding theme that more than any other appears to characterize the field of cognitive science. There are a number of ways of expressing this theme, for example, as the attempt to view intelligent behavior as consisting of the processing of information or as the attempt to view intelligence as the outcome of rule-governed activity (see RULE-BASED SYSTEMS). These characterizations express the same underlying idea. Computation, information processing, and rule-governed behavior all depend on the existence of physically instantiated codes or symbols that refer to or represent things and properties extrinsic to the behaving system. In all these cases the behavior of the systems in question (be they minds, computers, or social systems) are explained not in terms of intrinsic properties of the system itself but in terms of rules and processes that operate on representations of extrinsic things. Cognition, in other words, is explained in terms of regularities holding over semantically interpreted symbolic representations, just as the behavior of a computer evaluating a mathematical function is explained in terms of its having representations of mathematical expressions (eg, numerals) and in terms of the mathematical properties of the numbers these expressions represent. This is also analogous to explaining economic activity not in terms of the categories of natural science (eg, speaking of the physicochemical properties of money and goods) but in terms of the conventional meaning or symbolic value of these objects (eg, that they are taken to represent such abstractions as legal tender). Although in both economics and cognitive science the meaning-bearing objects (or the instantiation of the symbols) are physical, it is only by referring to their symbolic or referential character that we can explain the observed regularities in the resulting behavior.

There has been some misunderstanding of the significance of the assumption that cognition is explained in terms of regularities. For example, some people have suggested that this is no different from any other science, because all scientific theories deal with representations (eg, mathematical symbols that designate certain objects or properties). Hence simulations involving such theories (eg, simulations of planetary motions) are sometimes thought to be no different in principle from simulations of cognition. But the difference in two types of simulation is in fact fundamental, because in the case of cognition, the claim is that the organism being modeled, not just the theorist, actually manipulates physical tokens of the symbols, a claim that clearly has no parallel in physics unless the physicist is being modeled!

This representation thesis, sometimes referred to in philosophy as the representational theory of mind (Fodor, 1981) and in cognitive science as the physical-symbol system hypothesis (Newell, 1980, 1982) is one of the cornerstones of the discipline of cognitive science and is one of the features that links it in a fundamental way to AI. The intellectual and philosophical underpinnings of these two fields are now so closely linked that the distinction between them remains mostly at the pragmatic level, resting on such things as how big a role actual computer programs play and how technical are the immediate applications of the research. Some people expect that as cognitive scientists become better trained in computer science, and as AI begins to tackle the harder problem of what makes general intelligence possible, the distinction between the fields will fade. Similarly, the philosophy of mind is being influenced more and more by developments in AI and might be expected to play a more central role in clarifying the difficult conceptual issues that face both empirical and theoretical studies of intelligence.

BIBLIOGRAPHY

J. R. Anderson, "Acquisition of Cognitive Skill," *Psychol. Rev.* **89,** 369–406 (1982).

J. A. Anderson and G. E. Hinton, "Models of Information Processing in the Brain," in G. E. Hinton and J. A. Anderson, eds., *Parallel Models of Associative Memory*, Lawrence Erlbaum, Hillsdale, N.J., 1981, pp. 9–48.

R. C. Berwick and A. S. Weinberg, *The Grammatical Basis of Linguistic Performance: Language Use and Acquisition*, MIT Press, Cambridge, Mass., 1984.

M. Brady, *Artif. Intell.* (Special vol.) **17**(1–3), (Aug. 1981).

J. S. Brown and K. Van Lehn, "Repair Theory: A Generative Theory of Bugs in Procedural Skills," *Cogn. Sci.* **4,** 379–426 (1980).

S. K. Card, T. P. Moran, and A. Newell, *The Psychology of Human-Computer Interactions*, Lawrence Erlbaum, Hillsdale, N.J., 1983.

W. E. Cooper, *Cognitive Aspects of Skilled Typewriting*, Springer-Verlag, New York, 1983.

R. Cummins, *The Nature of Psychological Explanation*, MIT Press, Bradford Books, Cambridge, Mass., 1983.

W. Demopoulos and A. Marras, *Language, Learning and Concept Acquisition: Foundational Issues*, Ablex, Norwood, N.J., 1985.

D. Dennett, *Brainstorms*, MIT Press, Bradford Books, Cambridge, Mass., 1979.

K. Deutsch, "The Nerves of Government," *Gen. Syst. Yearbk.* **21,** 125–176 (1963).

G. W. Ernst and A. Newell, *GPS: A Case Study in Generality and Problem Solving*, Academic Press, New York, 1969.

J. Fodor, *Representations*, MIT Press, a Bradford Book, Cambridge, Mass., 1981.

J. Fodor, *The Modularity of Mind: An Essay on Faculty Psychology*, MIT Press, a Bradford Book, Cambridge, Mass., 1983.

J. W. Forrester, *World Dynamics*, Wright-Allen, Cambridge, Mass., 1971.

D. Gentner and A. L. Stevens, *Mental Models*, Lawrence Erlbaum, Hillsdale, N.J., 1982.

J. Haugeland, *Mind Design*, MIT Press, Bradford Books, Cambridge, Mass., 1981.

J. R. Hobbs and R. C. Moore, *Formal Theories of the Commonsense World*, Ablex, Norwood, N.J., 1984.

P. N. Johnson-Laird, *Mental Models*, Harvard University Press, Cambridge, Mass., 1983.

S. M. Kosslyn, *Image and Mind*, Harvard University Press, Cambridge, Mass., 1980.

M. Marcus, *A Theory of Syntactic Recognition for Natural Language*, MIT Press, Cambridge, Mass., 1979.

D. Marr, "Approaches to Biological Information Processing," *Science* **190,** 875–876 (1975).

D. Marr, *Vision*, W. H. Freeman, San Francisco, 1982.

R. S. Michalski, J. G. Carbonel, and T. M. Mitchell, *Machine Learning: An Artificial Intelligence Approach*, Vol. 2, Tioga Press, Palo Alto, Calif., 1986.

A. Newell, "Remarks on the Relationship between Artificial Intelligence and Cognitive Psychology," in R. Banerji and M. D. Mesarovic, eds., *Theoretical Approaches to Non-Numerical Problem Solving*, Springer-Verlag, New York, 1970.

A. Newell, "Artificial Intelligence and the Concept of Mind," in R. C. Schank and K. Colby, eds., *Computer Models of Thought and Language*, W. H. Freeman, San Francisco, 1973.

A. Newell, "Physical Symbol Systems," *Cogn. Sci.* **4,** 135–183 (1980).

A. Newell, "The Knowledge Level," *Artif. Intell.* **18,** 87–127 (1982).

Z. Pylyshyn, "Validating Computational Models: A Critique of Anderson's Indeterminacy of Representation Claim," *Psychol. Rev.* **86**(4), 383–394 (1979).

Z. Pylyshyn, "Complexity and the Study of Human and Machine Intelligence," in J. Haugeland, ed., *Mind Design*, MIT Press, Cambridge, Mass., 1980.

Z. Pylyshyn, "The Imagery Debate: Analogue Media versus Tacit Knowledge," *Psychol. Rev.* **88,** 16–45 (1981).

Z. Pylyshyn, "Information Science: Its Roots and Relations as Viewed from the Perspective of Cognitive Science," in F. Machlup and U. Mansfield, eds., *The Study of Information: Interdisciplinary Messages*, Wiley, New York, 1983, pp. 63–80.

Z. Pylyshyn, *Computation and Cognition: Toward a Foundation for Cognitive Science*, MIT Press, Cambridge, Mass., 1984.

R. C. Schank and R. P. Abelson, *Scripts, Plans, Goals and Understanding*, Lawrence Erlbaum, Hillsdale, N.J., 1977.

M. S. Seidenberg and M. K. Tanenhaus, "Modularity and Lexical Access," in I. Gopnik and Myrna Gopnik, eds., *From Models to Modules: Studies in Cognitive Science*, Ablex, Norwood, N.J., 1985.

H. A. Simon, *The Sciences of the Artificial*, MIT Press, Cambridge, Mass., 1969.

A. Sloman, *The Computer Revolution in Philosophy: Philosophy, Science, and the Models of Mind*, Humanities Press, New York, 1968.

D. Swinney, "Lexical Access during Sentence Comprehension: (Re)consideration of Context Effects," *J. Verb. Learn. Berb. Behav.* **18,** 645–660 (1979).

K. Wexler and P. Cullicover, *Formal Principles of Language Acquisition*, MIT Press, Cambridge, Mass., 1980.

S. Zucker, A Rosenfeld, and L. Davis, "General Purpose Models: Expectations about the Unexpected," in *Proceedings of the Fourth IJCAI*, Tbilisi, USSR, Morgan-Kaufmann, San Mateo, Calif., 1975.

Z. W. Pylyshyn
University of Western Ontario

ENCYCLOPEDIA OF ARTIFICIAL INTELLIGENCE

This extensively revised and expanded *Second Edition* of the *Encyclopedia of Artificial Intelligence* defines the discipline by bringing together the core of knowledge from all fields encompassed by AI. It covers the latest developments in current AI topics such as neural networks, fuzzy logic, machine vision, natural language generation, and many more. Includes:

- Over 450 articles—all entries written expressly for the *Encyclopedia*
- Over 5,000 literature references; 454 illustrations and color photographs
- Over 50% new and revised material
- Exemplary indexing and cross-referencing for easy, complete information access to all topics

Praise for the *First Edition* . . .

"The *Encyclopedia* is a wonder of clarity and scope: surprisingly easy to read . . . the clarity is an especially pleasant surprise, considering the articles were all written by AI experts . . . It's a treasure house of easily accessible knowledge."
—*Language Technology*

"Excellent bibliographies are attached to most of the articles, and diagrams and sketches are clear and helpful. The indexing and cross-indexing are exemplary. As the editor points out, the reader will be led by the extensive cross-references to almost every other article . . ."
—*Artificial Intelligence Reporter*

"The *Encyclopedia* is a first-class piece of work that will be an indispensable part of any AI library."
—*Computing Reviews*

". . . A tour de force . . . a truly fantastic encyclopedia which no one in the field of artificial intelligence can afford to be without."
—*Systems Research & Information*

WILEY-INTERSCIENCE
John Wiley & Sons, Inc.
Professional, Reference and Trade Group
605 Third Avenue, New York, N.Y. 10158-0012
New York • Chichester • Brisbane • Toronto • Singapore
ISBN 0 471-50307-X (Two-Volume Set)
ISBN 0 471-50305-3 (Vol. 1)
ISBN 0 471-50306-1 (Vol. 2)

ISBN 0-471-50305-3
90000
9 780471 503057