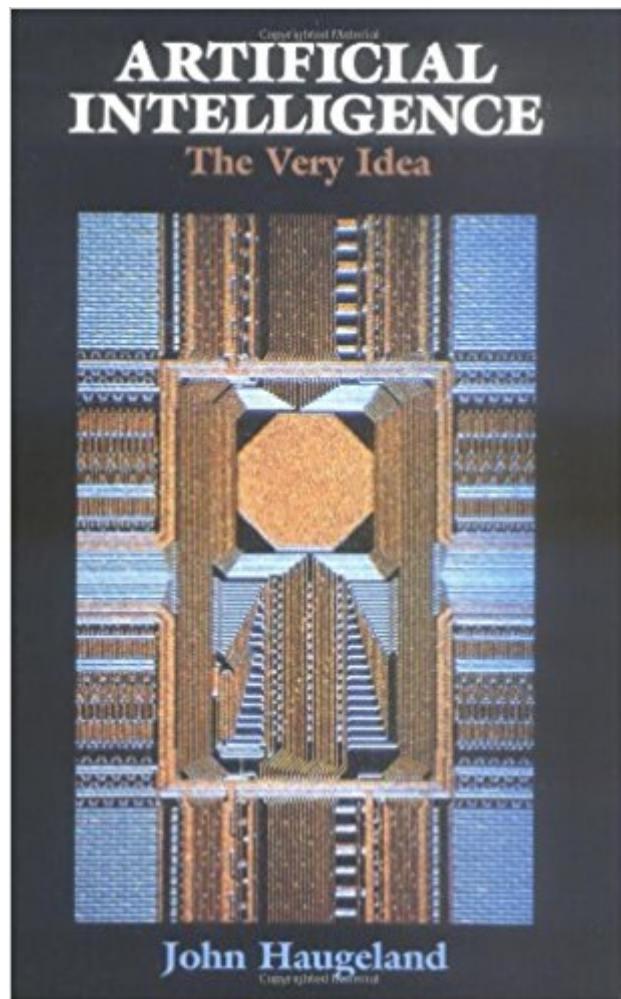


The book was found

Artificial Intelligence: The Very Idea



Synopsis

"Machines who think -- how utterly preposterous," huff beleaguered humanists, defending their dwindling turf. "Artificial Intelligence -- it's here and about to surpass our own," crow techno-visionaries, proclaiming dominion. It's so simple and obvious, each side maintains, only a fanatic could disagree. Deciding where the truth lies between these two extremes is the main purpose of John Haugeland's marvelously lucid and witty book on what artificial intelligence is all about. Although presented entirely in non-technical terms, it neither oversimplifies the science nor evades the fundamental philosophical issues. Far from ducking the really hard questions, it takes them on, one by one. Artificial intelligence, Haugeland notes, is based on a very good idea, which might well be right, and just as well might not. That idea, the idea that human thinking and machine computing are "radically the same," provides the central theme for his illuminating and provocative book about this exciting new field. After a brief but revealing digression in intellectual history, Haugeland systematically tackles such basic questions as: What is a computer really? How can a physical object "mean" anything? What are the options for computational organization? and What structures have been proposed and tried as actual scientific models for intelligence? In a concluding chapter he takes up several outstanding problems and puzzles -- including intelligence in action, imagery, feelings and personality -- and their enigmatic prospects for solution.

Book Information

Series: MIT Press

Paperback: 299 pages

Publisher: A Bradford Book; Reprint edition (January 6, 1989)

Language: English

ISBN-10: 0262580950

ISBN-13: 978-0262580953

Product Dimensions: 5.9 x 0.6 x 9 inches

Shipping Weight: 14.4 ounces (View shipping rates and policies)

Average Customer Review: 4.2 out of 5 stars 5 customer reviews

Best Sellers Rank: #337,685 in Books (See Top 100 in Books) #103 in Books > Textbooks > Computer Science > Artificial Intelligence #259 in Books > Computers & Technology > Computer Science > AI & Machine Learning > Intelligence & Semantics #462 in Books > Textbooks > Social Sciences > Psychology > Cognitive Psychology

Customer Reviews

This interesting book stands out from the many other new titles on the topic of artificial intelligence. Its philosophical treatment, which includes coverage of automatic formal systems, semantics, and the development of the various theories of thinking, presents an approach taken by few other AI works. The author ties together this philosophical treatment with clear explanations of how computer-based AI efforts operate and what this might hold for us in terms of future potential. He raises some thought-provoking questions and freely admits that the answers are not as clear-cut as some computer experts would have us think. This book should appeal to a wide audience. Hilary D. Burton, Livermore National Labs., Livermore, Cal. Copyright 1986 Reed Business Information, Inc. --This text refers to an out of print or unavailable edition of this title.

A delightfully well written book highlighting many of the deep and important issues forming the foundation of artificial intelligence and, in fact, of all cognitive science. (Contemporary Psychology)[An] amusing and wide ranging study... As well as containing concise and sharp definitions of fundamental issues in the philosophy of logic and meaning, the book has excellent summaries of basic computer architectures and hot topics in AI research. (Times Literary Supplement)

John Haugeland's book begins with an hypothesis: human thought is at bottom "rational symbol manipulation." This is a contentious thesis but Haugeland does not present it dogmatically. Haugeland's book presents a narrative history of AI and the attempts that have been made to build realistic computers that are capable of mimicking human thought by manipulating symbols. The book is broken up into topics but Haugeland also manages to present a lot of the history and context which is good for neophytes like me. John Haugeland's book has a number of virtues that make it a good introduction to classical AI. First, its accessibility. I have a strong background in philosophy, but I have almost no background in philosophy of mind, computer science, or cognitive science. Even without that background I found the ideas presented in this book to be both lucid and powerful. The chapter that examines computer architecture was the most challenging but Haugeland still manages to present the ideas clearly. Second, Haugeland is honest about the serious difficulties that classical AI runs into when trying to produce realistic models of human thought. However, as I realized very quickly reading this book, even the failures of classical AI are illuminating. We learn a great deal about human thought by examining the ways in which computers fail to realistically mimic our thought. By seeing what computers are (at the present moment) incapable of doing we get some insight into what we are capable of doing. I think human thought

progresses by following paths to their natural conclusion and, even when we reach dead ends, we have learned something about why that path was a dead end. So, I think it is great that there are people taking the basic hypothesis that human thought is rational symbol manipulation as far as it will go. Even if the version of AI presented in this book turns out to be a dead end we will know more about why it is a dead end if we pursue it to the end. If we simply dismiss it on a priori grounds we gain nothing. Third, Haugeland himself is not dogmatic one way or the other, which is extremely refreshing. He clearly finds the prospect of AI exciting but he does not consider it inevitable. He sees the problems and challenges clearly and he does not hide them from the reader or adopt a defensive posture. It is easy to dismiss the idea of AI, but the easy dismissal of AI is generally based on our intuitive understanding of how our thought functions, and our intuitions are often wrong. One of the great things about AI research is that it often presents us with really startling, counter-intuitive hypotheses about how our minds function. And finally, Haugeland never pretends we know more than we do and he leaves the reader with lots of unsolved problems begging for thought and future research (which is a good thing). All in all, if you are looking for a good, accessible, exciting introduction to classical AI, Haugeland's book is a great place to start.

The state of play in artificial intelligence has changed considerably since this book was written in the mid-80's, but that doesn't mean that this book has become outdated. Far from it. The subsequent changes in artificial intelligence have merely served to highlight the fact that the value of this book is not in its discussion of 80's technology. The value of this book is that it articulates, motivates, and then critiques an idea at the core of philosophical thinking about mindedness over the past 400 years: namely, the idea that cognition consists in the manipulation of inner representations. Although Haugeland himself went on to offer a radically different account of mindedness, according to which minds are essentially embodied and embedded, the value of this book does not derive from its articulation of an original view. Rather, the value of it derives from the rigor with which Haugeland motivates the attractions of a philosophical view that he himself does not accept. Buy this book if you want to understand the underlying philosophical motivations for the widespread view that thoughts are inner representations. Buy Haugeland's collection of essays, entitled "Having Thought", if you want to engage with Haugeland's own alternative view of the mind.

Don't judge this book by its cover—or at least by its title. Haugeland's *Artificial Intelligence: The Very Idea* does not adequately serve as a general introduction to the conceptual underpinnings and philosophical background of the quest to create an artificial mind. Rather, it focuses on one specific

approach to how natural and man-made thought works: "thinking...essentially is rational manipulation of mental symbols." (p. 4) Haugeland plows forward with this as his core assumption, barely noting that some AI researchers see thought from a very different perspective (for example, the connectionists) and others find the whole enterprise fraught with theoretical difficulty (such as Dreyfus). So Haugeland's story is that of a particular theory of mind that held predominance for several decades (what the author himself dubs "good, old-fashioned artificial intelligence" or "GOFAI", p. 112) but is now gradually being superceded. His introduction to this story concludes with a description of the Turing test and a justification for its use, and a brief statement of the efficacy of describing a system in different-even contradictory-ways through different "organizational levels". (p. 9) Of all the ideas presented in the book, this last one has the greatest promise for applicability beyond GOFAI. Chapter 1, "The Saga of the Modern Mind", is a condensed bit of intellectual history. Haugeland introduces the philosophical children of the Copernican revolution-Hobbes, Descartes, and Hume-and the ways they grappled with understanding the world of the mental with the ideas that had proven so effective in the physical sciences. We soon encounter the "paradox of mechanical reason": if reason is the meaningful manipulation of symbols, and meanings are not physical entities, then how can machines manipulate them? (p. 39) Chapter 2 serves as an extended definition of "Automatic Formal Systems", that is, computers. This material is the most challenging in the text, but the important concepts (formal games, digital systems, medium independence, etc.), are well-described, except for finite playability. The students I tutored through this work found it impossible to determine just what point was being made, and so did I. How does one assign meanings-connections to the "real", outside world-to the symbols that a computer manipulates? This question is taken up in Chapter 3, "Semantics"-and answered, it seems, by sleight-of-hand. Haugeland gives to this the name "the formalist's motto": "if you take care of the syntax, the semantics will take care of itself". (p. 106) Neither I nor my students found this simple resolution at all satisfying. In every example of a formal game that the author presents, whatever semantic interpretation it has is provided from outside the system. Chapter 4, "Computer Architecture", charts the milestones of computing. It begins with the analytical engine, and lauds Babbage's single-handed invention of programming without noting, however, that a human mind does not resemble the tabula rasa of a computer's memory bank. Moving quickly to the twentieth century, we get insightful descriptions of Turing machines, von Neumann machines (which turn out to be the kind of computer we are accustomed to), the mind-bending tree-structured LISP machines, and Newell's pragmatic production machines. Chapter 5, "Real Machines", might be better titled "Real Problems". Haugeland presents some of the brick walls that AI research has run into. These

can be grouped into the phenomenon of the combinatorial explosion: in order to interact with the real world in a manner that demonstrates "common sense", an AI must have access to an impossibly large store of information (while accessing what it needs in due time), and be able to consider an equally impossibly large set of potential courses of action. (p. 178) Methods to restrict what the AI has to consider, such as the focus on "micro-worlds", result in a system with no sense. Haugeland acknowledges these problems, and offers nothing but hope in scientific and technological progress to answer them. Chapter 6, "Real People", develops means by which the sense that humans exhibit, and machines are far from realizing. Dennett's intentional stances and Grice's conversational implicatures are intelligent-if partial-characterizations of perspicuous reasoning. They are, however, frustratingly slippery for computer programmers, so it's not surprising that Haugeland, with some exasperation, groups them together under the "nonasinity canon": "An enduring system makes sense to the extent that, as understood, it isn't making [a rear] of itself." (p. 219) I feel that, if a reader has followed the author this far, then he or she deserves better than this. Yet Haugeland and his colleagues are bound to feel frustration. Computers are electromechanical in nature, while humans are neurochemical. Computers can engage in numerical calculation with speed and precision, while most people find mathematics to be their most difficult school subject. Computers are tools that we devised to assist us. Human behavior was forged in the four-billion cauldron of evolution, and psychologists have barely begun to sort out the seething stew of vestigial loves, hates, and motivations that shape our behavior. And honest cognitive science will admit that humans and supercomputers are each masters of two separate, very different worlds. At the end, Haugeland finally admits this possibility-without contemplating the alternatives to the computation theory of might that this possibility demands.

This is a great exposition of the fundamental notions involved in the philosophy of AI. While at first look may appear like a good undergraduate read, it is, in fact, quite subtle and deep in most of the material it touches. Great scholarship.

[Download to continue reading...](#)

Readings in Medical Artificial Intelligence. The First Decade (Addison-Wesley Series in Artificial Intelligence) Artificial Intelligence: The Very Idea Emotional Intelligence: Why You're Smarter But They Are More Successful(Emotional intelligence leadership,Emotional Quotient,emotional intelligence depression,emotional intelligence workbook) Emotional Intelligence: 3 Manuscripts - Emotional Intelligence Definitive Guide, Mastery, Complete Step by Step Guide (Social Engineering, Leadership, ... (Emotional Intelligence Series Book 4) Life 3.0: Being Human in the Age of Artificial

Intelligence Artificial Intelligence for Games Robots and Artificial Intelligence (Technology Behind) The Fourth Transformation: How Augmented Reality & Artificial Intelligence Will Change Everything Humans Need Not Apply: A Guide to Wealth and Work in the Age of Artificial Intelligence Artificial Intelligence: What Everyone Needs to Know Economy Monitor Guide to Smart Contracts: Blockchain Examples (Artificial Intelligence, Law and Finance) The Most Human Human: What Artificial Intelligence Teaches Us About Being Alive Applications of Artificial Intelligence for Decision-Making: Multi-Strategy Reasoning Under Uncertainty Essentials of Game Theory: A Concise, Multidisciplinary Introduction (Synthesis Lectures on Artificial Intelligence and Machine Learning) Forbidden Gates: How Genetics, Robotics, Artificial Intelligence, Synthetic Biology, Nanotechnology, & Human Enhancement Herald The Dawn Of Techno-Dimensional Spiritual Warfare Artificial Intelligence in Label-free Microscopy: Biological Cell Classification by Time Stretch Artificial Intelligence In Medicine (Aaas Selected Symposium) The Technological Elixir: Invoking Artificial Intelligence & Understanding the Arcana of Mysticism Singularity Now! The Artificial Intelligence Timeline: Automation, robotics, tech titans, and their brave new world (September 2017) TI vs AI: Targeted Individual vs. Artificial Intelligence

[Contact Us](#)

[DMCA](#)

[Privacy](#)

[FAQ & Help](#)