AI Magazine Volume 12 Number 4 (1991) (© AAAI)

cerned with the foundations of natural language processing, but AI practitioners building systems today will find it of little appeal. Because it assumes so much previous knowledge, the book will not be useful to the casual reader. One would be at a disadvantage without a reasonable familiarity with predicate calculus and modal logic, AI planning formalisms, and the work of Perrault and Allen on interpreting speech acts (for example, Allen and Perrault [1980]; Perrault and Allen [1980]). Accordingly, the reader of this review should be warned that my point of view is that of a researcher (specifically, an academic researcher) rather than a system builder; your mileage might vary.

No review of this book would be complete without some mention of the commentaries, critical pieces written by other workshop participants that follow groups of related papers. Each commentator did an excellent job. The inclusion of these well-considered short pieces helps focus the reader's attention on important features of the related papers, giving him/her a feeling of participation in a fascinating discussion. The editors deserve congratulations for their fine work in editing and arranging. *Intentions in Communication* is one of the best-edited collections I have had the privilege to read. This point is particularly laudable in light of the book's origin in a workshop, which often makes for slapdash publications.

In fact, my only serious argument with the book is that it might have been improved by the editors being even more of a presence. A more substantial introductory chapter that gave more background and, perhaps, even a glossary would have opened the book to more readers. This need is particularly true with the casual reader, who must grapple with a dialect that consists of the jargons of philosophy, AI, and linguistics together. An expanded introduction and the inclusion of a glossary might also have spared the reader from wading through four or five fragmentary introductions to the work of Austin and Searle.

I borrow the words of Abraham Lincoln to conclude: "People who like this sort of thing will find this is the sort of thing they like." Not for all readers, *Intentions in Communication* is essential for those interested in foundational issues of natural language processing. Certainly, no research library should be without it.

### References


Allen, J. F., and Perrault, C. R. 1980. Analyzing Intention in Utterances. *Artificial Intelligence* 15(3): 143–178.

Austin, J. L. 1962. *How to Do Things with Words*. London: Oxford University Press.

Cohen, P. R., and Perrault, C. R. 1979. Elements of a Plan-Based Theory of Speech Acts. *Cognitive Science* 3:177–212.

Grice, P. 1975. Logic and Conversation. In *Syntax and Semantics*, eds. P. Cole and J. Morgan, 41–58. New York: Academic.

Hanks, S., and McDermott, D. 1986. Default Reasoning, Nonmonotonic Logics, and the Frame Problem. In Proceedings of the Fifth National Conference on Artificial Intelligence, 328–333. Menlo Park, Calif.: American Association for Artificial Intelligence.

Kautz, H. 1987. A Formal Theory of Plan Recognition, Technical Report, TR 215, Dept. of Computer Science, Univ. of Rochester.

Kautz, H., and Allen, J. 1986. Generalized Plan Recognition. In Proceedings of the Fifth National Conference on Artificial Intelligence, 32–38. Menlo Park, Calif.: American Association for Artificial Intelligence.

McCarthy, J. 1980. Circumscription—A Form of Nonmonotonic Reasoning. *Artificial Intelligence* 13:27–39, 171–172.

Perrault, C. R., and Allen, J. 1980. A Plan-Based Analysis of Indirect Speech Acts. *American Journal of Computational Linguistics* 6(3–4): 167–182.

Reiter, R. 1980. A Logic for Default Reasoning. *Artificial Intelligence* 13:81–132.

Searle, J. 1969. *Speech Acts: An Essay in the Philosophy of Language*. London: Cambridge University Press.


**Robert P. Goldman** is an assistant professor of computer science at Tulane University. He received his Ph.D. in computer science from Brown University for work on the application of probability theory to natural language processing.

# On Being a Machine

*Gabriel Valiente*

*On Being a Machine, Volume 1: Formal Aspects of Artificial Intelligence,* A. Narayanan, Ellis Horwood Limited, Chichester, England, 1988, 200 pages, hardback, US$39.95, ISBN 0-85312-957-6.

A great debate concerning the possibility for machine intelligence began with the advent of computing. Roughly stated, in one corner, we find AI researchers and practitioners developing computational models that exhibit an ever-increasing degree of intelligence, and in the opposite corner, we find formal theoreticians, philosophers, and psychologists arguing about the fundamental capabilities and limitations of machines. This debate seems to have no end because although there are strong arguments that shake the conceptual foundations of AI, counterarguments are as strong, and none presents decisive, irrefutable evidence of the basic capabilities or limitations of machines.

Ajit Narayanan's book *On Being a Machine, Volume 1: Formal Aspects of Artificial Intelligence* sheds new light on these issues by providing a formal analysis of the main arguments and counterarguments of AI proponents and critics. The book gives an introduction to this subject from a formal basis that is suited for a wide audience, including computer scientists, AI researchers and practitioners, formal theoreticians, philosophers, and psychologists. A description of the main formal aspects of AI is also provided, and the possibility for machine intelligence is analyzed from this formal standpoint.

The essence of the author's perspective is that "AI, despite considerable advances in its techniques, tools and applications, has not developed significantly as far as its theoretical and philosophical aspects are concerned, because from the very beginning AI has been miscategorized by theoreticians, philosophers, and even AI researchers" (p. 9). In addition, the cause of the miscategorization is rooted in the acceptance that all formal limitations that apply to computer science and philosophy apply to AI as well.

The basis for analyzing these formal limitations is Turing's imitation game and the objections to the game that Turing himself formulated. Although other arguments are covered as well (Minsky's proof of unsolvability of the halting problem, Searle's Chinese room argument, and Lucas's version of the mathematical objection), many important arguments to the possibility for machine intelligence are not mentioned at all, among them computational complexity limits (that is, time and space

complexity of algorithms, complexity inherent to the task of specifying the deterministic or nondeterministic machine, and complexity of electric or logical circuits), physical limits of computing (that is, computation in the physical world requires the expenditure of energy, communication in space, and the passage of time), and limits of conceptualization (that is, finite, discrete concepts can never form a perfect model of a continuous world; the only things that can accurately be represented in concepts are manmade structures that once originated as concepts in some person's mind or systematic domains of distinctions created through the use of language) (Sowa 1984; White 1988). Furthermore, the main conceptual foundations of AI—namely, the knowledge representation hypothesis of Brian Smith (1982) and the physical symbol system hypothesis of Allen Newell (1980)—are not discussed at all. These hypotheses have been considered fundamental cornerstones of AI research, but they are now being questioned as posing strong limitations on AI (Dahlbäck 1989; Dreyfus 1972; Winograd and Flores 1986).

Given this perspective, the author concludes that AI's essential methodology is a continuous attempt to overcome the formal constraints of computer science and philosophy without sacrificing rigor. Although I liked the author's perspective, and I wholly agree with his main conclusion, both are just stated in the preface, and no further reference to them is given.

Let's get a feeling of what this first volume is really about. The organization of the text is clear and straightforward. This first volume is mainly expository in nature and comprises a detailed discussion of AI's formal constraints as a starting point to the discussion of AI's essential methodology that is promised for the second volume. The question of the possibility of machine intelligence is introduced in the first chapter through a detailed exposition on Turing's imitation game and an informal statement of its original objections. Turing's reformulation of the question of the possibility of machine intelligence into the imitation game is discussed after a brief exposition of two doctrines in the philosophy of mind—naive dualism and naive logical behaviorism—that provide the necessary background for an interpretation of the reformulation. Each of the objections is then analyzed from a formal standpoint because the relevant elements of formal theory are introduced in subsequent chapters.

Theoretical computer science is dealt with in chapter 2. The chapter contains a brief description of function and automata theory as a basis for the detailed analysis of Lady Lovelace's objection. Despite the introductory character of the chapter, the omission of the theory of (primitive) recursive functions is surprising because Turing's and Church's theses on the limitations of computability, as well as Gödel's theorems, are strongly based on this theory (see, for example, Delong [1970]).

Formal philosophy is addressed in two chapters, one devoted to logic and another to semantics. Both propositional and predicate logic are covered in chapter 3 in addition to formal theories and systems, theorem proving, and logic-based knowledge representation. Gödel's theorems are exposed with clarity, although oversimplified, and the Mathematical objection in both its original and Lucas's version is thoroughly discussed. The presentation is far from complete, particularly with respect to theorem proving and logic-based knowledge representation (a more thorough treatment can be found, among others, in Genesereth and Nilsson [1987]), but is detailed enough to understand the metatheoretic view of logic systems.

Chapter 4 is much more comprehensive. It provides a detailed treatment of truth-conditional, model, and possible-world semantics. With the background in tensed modal logic given by the exposition of possible-world semantics, the distinction between the possibility and the necessity for machine intelligence is depicted with clarity. The argument from informality of behavior is also discussed.

In summary, this first volume provides a readable introduction to formal foundations of AI and gives a comprehensive analysis of the possibility for machine intelligence from this formal standpoint. Its significance is that it introduces a wide audience to the main issues surrounding thinking machines as well as provides a formal analysis of informal arguments for and against the possibility of machine intelligence. I highly recommend the book to anyone interested in the AI debate.

The second volume promises to draw on a characterization of AI's essential methodology as continuous attempts to overcome the formal constraints of computer science and philosophy by augmenting appropriate formal theories with nonformal yet rigorous models and approaches. It will also cover recent developments in neurocomputing. I hope to see my criticisms dissipate after reading the second volume.

### References


Dahlbäck, N. 1989. A Symbol Is Not a Symbol. In Proceedings of the Eleventh International Joint Conference on Artificial Intelligence, 8–14. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.

Delong, H. 1970. *A Profile of Mathematical Logic*. Reading, Mass.: Addison-Wesley.

Dreyfus, H. L. 1972. *What Computers Can't Do*. New York: Harper and Row.

Genesereth, M. R., and Nilsson, N. 1987. *Logical Foundations of Artificial Intelligence*. San Mateo, Calif.: Morgan Kaufmann.

Newell, A. 1980. Physical Symbol Systems. *Cognitive Science* 4:135–183.

Smith, B. 1982. Reflection and Semantics in a Procedural Language, Technical Report, MIT/LCS/TR-272, Massachusetts Institute of Technology.

Sowa, J. F. 1984. *Conceptual Structures: Information Processing in Man and Machine*. Reading, Mass.: Addison-Wesley.

White, I. 1988. The Limits and Capabilities of Machines: A Review. *IEEE Transactions on Systems, Man, and Cybernetics* 18(6): 917–938.

Winograd, T., and Flores, F. 1986. *Understanding Computers and Cognition: A New Foundation for Design*. Norwood, N.J.: Ablex.



**Gabriel Valiente** is a research associate in the Center of Advanced Studies of Blanes, Spain, and a graduate student at the Autonomous University of Barcelona. His research interests include the formal development and validation of expert systems.


## The Cognitive Structure of Emotions

*Roland J. Zito-Wolf*

*The Cognitive Structure of Emotions*, Andrew Ortony, Gerald L. Clore, and Allan Collins, Cambridge University Press, Cambridge, United Kingdom, 207 pages, $14.95, ISBN 0-521-38664-0.