# Introduction to Artificial Intelligence

Eugene Charniak

Drew McDermott

# Introduction to Artificial Intelligence

**Eugene Charniak**
Brown University

**Drew McDermott**
Yale University

ADDISON-WESLEY PUBLISHING COMPANY

Reading, Massachusetts ● Menlo Park, California ● Don Mills, Ontario
Wokingham, England ● Amsterdam ● Sydney ● Singapore ● Tokyo ● Mexico City
Bogotá ● Santiago ● San Juan

Mark S. Dalton/Sponsoring Editor

Hugh J. Crawford/Manufacturing Supervisor
Martha K. Morong/Production Manager
Dick Morton/Art Editor
Margaret Pinette/Production Editor
Patti O'Hare Williams/Text and Cover Designer
Kenneth Wilson/Illustrator
William J. Yskamp/Copy Editor

**This book is in the Addison-Wesley Series in Computer Science.**
Michael A. Harrison/Consulting Editor

***Credits***

Fig. 1.1: © Tom McHugh. Fig. 1.2: Adapted from Woods, William A., "An experimental parsing system for transition network grammars," in R. Rustin, *Natural Language Processing*. New York: Algorithmics Press, Inc., 1972. Fig. 1.3: Reprinted by permission of Westview Press from *Artificial Intelligence in Medicine*, edited by Peter Szolovits. Copyright © 1982 by the American Association for the Advancement of Science. Fig. 1.4: Reprint permission from *AI Magazine*. Vol. 2, No. 2, (1981) published by the American Association for Artificial Intelligence. Fig. 1.6: (right) Leonard Lee Rue; (left) Ylla. © 1978. Fig. 3.34: From Stevens, K.A., "Computation of locally parallel structure," *Biol. Cybernetics* pp. 19–28 (1978). Fig. 3.40: From *Vision* by David Marr. W.H. Freeman and Company. Copyright © 1982. Figs. 3.45–3.50: From Schatz, Bruce R., "The computation of immediate texture discrimination." Memo 426. MIT Artificial Intelligence Lab (1977). Fig. 3.57: From Barrow, H.G., and Tenenbaum, J.M., "Recovering intrinsic scene characteristics from images," in Hanson, A., and Riseman, E. (eds.), *Computer Vision Systems*. New York: Academic Press (1978). Fig. 3.63: From Barrow, H.G., and Tenenbaum, J.M., "Recovering intrinsic scene characteristics from images," in Hanson, A., and Riseman, E. (eds.), *Computer Vision Systems*. New York: Academic Press (1978). Fig. 3.81: From Shepard, Roger N., and Metzler, Jacqueline, "Mental rotation of three-dimensional objects," *Science* 171:701–703 (1971). Copyright 1971 by the American Association for the Advancement of Science. Fig. 3.82: From Hinton, Geoffrey, "Some demonstrations of the effects of structural descriptions in mental imagery," *Cognitive Science*. 3(3): 231–251 (1979). Fig. 4.2: Courtesy of Professor Harvey Silverman, Department of Engineering, Brown University. Fig. 9.38: From Wilkins, David, "Using patterns and plans in chess," *Artificial Intelligence* 14(2):165–203 (1980). Fig. 9.40: RISK® game equipment used by permission of Parker Brothers © 1975, 1980. RISK is Parker Brothers' registered trademark for its world conquest board game equipment. Fig. 11.20: Lew Merrim/Monkmeyer Press Photo Service, Inc.

**Library of Congress Cataloging in Publication Data**

Charniak, Eugene.
Introduction to artificial intelligence.

Includes bibliographies and index.
1. Artificial intelligence. I. McDermott, Drew V.
II. Title.
Q335.C483 1984 001.53'5 84-14542
ISBN 0-201-11945-5



Reproduced by Addison-Wesley from camera-ready copy supplied and approved by the authors.



ABCDEFGHIJ-HA-898765

# Preface

While we have no evidence, we suspect that most professors write textbooks because they are going to teach a course for which none of the existing texts are suitable. Whether or not this is typically the case, it was certainly the reason for the book you are now holding. There were three things we wanted from a text on Artificial Intelligence (from now on "AI").

- It had to include those aspects of the field that we felt would prove to be enduring (for example, low-level vision research, which is given short shrift in most AI textbooks).
- It should present AI as a coherent body of ideas and methods, drawing on the work of the last twenty years, but translating it into a more uniform set of concepts, favoring clarity over journalistic accuracy.
- It should acquaint the student not only with the classic programs in the field, but the underlying theory as well: mathematical logic for knowledge representation and inference, linguistics for language comprehension, etc.

In brief, our belief was that AI is distinguished by its subject matter, not its history or methods. Its subject matter is the mind, considered as an information processing system.

When each of us became aware that the other was thinking along these lines, we decided to pool our efforts and write a textbook to reflect our view of the field. Four years later, here is the result.

Writing a textbook turned out to be a good way to develop humility about the shortcomings of other textbooks. For one thing, a general-purpose textbook must cover the significant results in *all* aspects of the field, including some its authors know little about. So we had to do some reading and writing we hadn't anticipated.

Second, we have been made aware just how hard it is to extract coherent and lasting theories from a field developing as rapidly as AI. While we have succeeded to some degree in finding some unifying concepts and notations, the book still has aspects of the usual AI "grab-bag." We managed to throw out a few dusty programming techniques, but there are many areas where there is no satisfactory theory as yet, so we have had to present the traditional, unsatisfactory one instead. We have felt free to "editorialize" at several points about the inadequacies of current theories.

On balance, however, we are pleased by our book. We have some strong opinions about what good AI is, and we were gratified about how much of it there is to report on, and how often our opinions meshed into a coherent framework for reporting on it. We hope the audience sees the same coherence we did.

While we wrote this book for introductory undergraduate AI courses, at the second to fourth year level, there is wide variety in the kinds of students who might take such a course. "Ideal" students would come to the study of AI with a rather long list of previous courses. From computer science they would have had introductory programming (including recursion), data structures, logic (including the first order predicate calculus), and a smattering of formal language theory. From linguistics, they would have had a course in generative grammar, and perhaps one in semantics as well. Naturally, it would help if they had the calculus, a first course in physics (so as to better understand the work in low-level vision) and a course in statistics (for the treatment of medical diagnosis).

Naturally, *no* student will have this panoply of prerequisites. Furthermore, since we wanted the book to be usable in a variety of classroom situations, we did not want to absolutely assume *any* of the prerequisites, except some programming experience. So, while a course in AI for cognitive scientists should have people with linguistics backgrounds, they may or may not know much about logic. On the other hand, while upper level undergraduates in computer science will have had most of the computer science background, they may not have had formal language theory, and, if our experience has been typical, their grasp of predicate calculus will be tenuous.

Our solution to this problem has been to try to explain everything. However, typically our explanations will be fast, and thus instructors would be well advised to give extra material on those areas in which their students will be unprepared. One exception is the differential and integral calculus, which, for obvious reasons, we did not want to summarize. Fortunately, the calculus is only needed in one section on vision, and it can be skipped if need be. We have also tried to separate more advanced material into optional subsections. Thus we hope that this book can find a home in many climates.

Everything in the book depends on Chapter 1, in which we introduce the idea of "internal representation" for facts and rules. The need to devise flexible and efficient representations occurs in almost all subfields of AI.

Chapter 2 is an introduction to the programming language Lisp and is optional since we have made the rest of the text pretty much independent of the

language. Thus it should be possible to teach the course using another symbol-manipulation language, like Prolog. If you do, let us know. Indeed, one could read the book and profit from it without touching a computer, but it wouldn't be much fun. Part of the attraction of AI is getting a machine to exhibit some "mental" faculty, and readers shouldn't deny themselves this experience.

Chapters 3 and 4 examine transduction of information into internal representation for sight and language, the two senses most intensively studied by AI researchers.

The central chapters of the book are 6 and 7, which deepen the ideas sketched in Chapter 1 on internal representation and reasoning. We recommend doing one of 3 and 4 before these chapters, to see where the internal representations come from. Chapter 5, on search, presents the notion of a goal tree, which is used in many of the subsequent chapters.

Chapters 8, 9, 10, and 11 cover four important topics, reasoning under uncertainty, robot planning, language understanding, and learning. For the most part they are independent of each other, except that the work in Chapter 10 on motivation analysis depends on Chapter 9.

In summary, the order to do the book in is this:

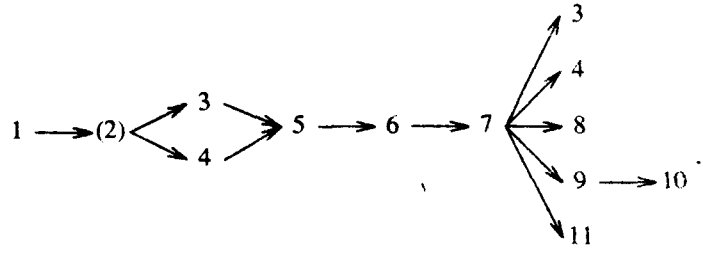


Chapter 2 is in parentheses because readers not interested in programming can skip it. Chapters 3 and 4 appear twice because they can be done in either place, although at least one of them ought to be done before 5. There are some local interchapter dependencies not shown in this diagram. These are indicated in the opening material at the front of each chapter.

Through out all of the chapters we have tried to give the reader some idea of the history of the field, as well as the intellectual filaments that connect AI to related disciplines. For this reason we have included various "boxes" covering topics which, while connected to the main body of the text, are nevertheless excursions. There are boxes on the history of AI, its relation to cognitive science, its relation to the rest of computer science, details of Lisp, etc. Often diversions are the most enjoyable parts of class lectures, and we hope these boxes give something of the same flavor.

Naturally, the most important way to make connections is through references, and we know from experience that many textbooks find their greatest use as reference works later. We have provided as many pointers into the AI literature as possible, but no effort has been made to find the earliest references to

some idea for the sake of giving credit. Rather we have tried to confine our references to those that would be useful to someone today, and that are published in accessible places, such as books, journals, and conference proceedings. We probably have not succeeded at this, but have come close enough to offend someone who, if this were a research paper, would have had his out-of-print, but seminal, technical report mentioned. We apologize on both counts.

*Acknowledgements*

This book benefited from the advice of many people. Early drafts were read and constructively criticized by James Allen, James Beisaw, Mark Burstein, Ernie Davis, Tom Dean, Michael Gavin, Fernando Gomez, Jim Hendler, Graeme Hirst, Lewis Johnson, Janet Kolodner, Michael Lebowitz, Stan Letovsky, Barbara Meier, David Miller, Ashfaq Munshi, Chris Riesbeck, Yoav Shoham, Jim Spohrer, Peter Wegner, and Steve Zucker. We got general advice and other help from Larry Birnbaum, Mike Brady, David Byrne, Greg Collins, Eric Gold, Matt Kaplan, Wendy Lehnert, Steve Reiss, and Chris Riesbeck. The brievity of our thanks should not belie the real impact all of these folks have had on the form and content of the book.

We would like to take this opportunity to note that the order of the names on the cover is alphabetical, and has no other significance.

E.C.
D.V.M.

# Contents

# 1

*Notes in this position at the beginning of chapters give the reader guidelines on what previous material is required for the chapter, and to what degree it is important that the chapter be read in sequence. Obviously, this chapter requires no previous material. All subsequent chapters make use of material covered here.*

# AI and Internal Representation

## 1.1 Artificial Intelligence and the World

As with many branches of computer science, it is unlikely that the average reader will come to this book with no idea of what Artificial Intelligence (AI) is, or what it is good for. At times it seems that one cannot read a newspaper or weekly magazine without some reference to the field, and the interest is certainly understandable. AI researchers are trying to create a computer which thinks. The very idea is intriguing, and reflection does not lessen the hold that such ideas have. What would the world be like if we had intelligent machines? What would the existence of such machines say about the nature of human beings and their relation to the world around them? Would college professors become obsolete?

We raise these questions not because this book will answer them; it will not. Ultimately these are questions about economics, psychology and, at the deepest level, the nature of human values. Rather we mention them to illustrate the ramifications of the science we will be describing in this book.

But even at a more mundane level of everyday economics, AI has tremendous implications. One subarea of AI is robotics. Figure 1.1 shows the use of robots to weld car bodies together. Many of the issues that arise in the construction of such robots are more properly in the domain of mechanical engineering than AI (e.g., can one create a motor which will allow a mechanical arm to move in certain ways?). Nevertheless questions like how to get a robot arm from one place to another without killing anybody are standard AI problems.

Another area of AI research is *natural language comprehension*. (We say "natural" language to distinguish it from computer languages.) Figure 1.2 shows examples of the kinds of question answered by a database system which

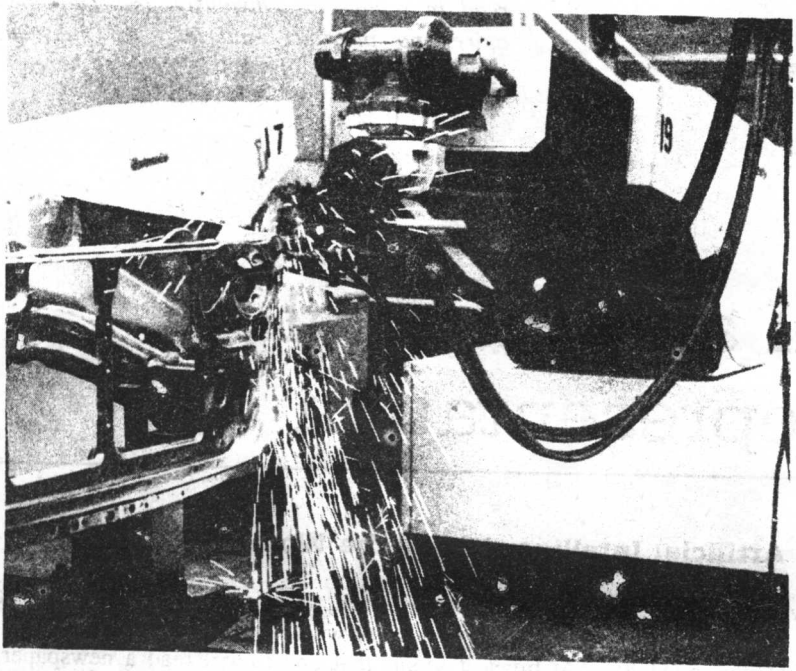

**Figure 1.1** Robot welding machines

```
** (How many samples contain chromite?)   ; The English query.
Interpretations:                           ; The translation into the
(printout                                  ; database query language.
 (number x6 / (seq samples) :              ; Asked to print out the
   (and T (contain x6 'chromite))))        ; number of things x6 that are
                                           ; samples containing chromite.
Executing: (3)                             ; There were three.
** (What are those samples?)               ; Second query.
Interpretations:                           ;
(for every x6 / (seq samples) :            ; For every sample x6 that
     (and T (contain x6 'chromite))        ; contains chromite
     (printout x6))                        ; print it out.
Executing: s10020                          ; The samples are s10020
           s10045                          ; s10045
           s10084                          ; and s10084.
```

**Figure 1.2** Examples of the Lunar database query system

has a *natural language front end*. That is, there is a program which translates the query from natural language to some database query language — an artificial language designed to be used by computer databases. The particular front end we show is a relatively old one, the *Lunar* system [Woods72]. This program answered questions about the rock samples which were brought back from the moon.

Other researchers have been trying to get machines to do medical diagnosis. Figure 1.3 shows some output from the *Caduceus* program [Miller82]. This is a program which, given the symptoms of the patient, will suggest other tests to be performed, and once it has enough information will suggest what disease or diseases the patient has. According to its creators this program knows about two-thirds of the diseases which go under the heading of *internal medicine* (medicine of the internal organs, as opposed to, say, the eyes or the bones). Admittedly, doctors are still better than the machine, but progress has been considerable.

One AI program that has replaced people in its domain of expertise is the *Xcon* program [McDermott81] (originally called R1). Xcon *configures* computers (and in particular, the Digital Equipment Corporation's Vax™ computer).

```
(Doctor)
Internist-1 consultation Sumex-aim ; Caduceus was originally named
version                            ; Internist. Sumex-aim is a group
                                   ; of researchers in AI medicine who
                                   ; share computational resources.
Please enter findings              ; The program first asks the
* sex male                         ; doctor to enter the facts about
* race white                       ; the patient.
* age 26 to 55                     ; There is a fixed vocabulary of
* alcoholism chronic               ; symptoms which must be followed.
* go                               ; This tells Caduceus to take over.
Disregarding:                      ; Caduceus finds a set of diseases
  exposure to rabbits              ; of which it suspects the patient
  leg weakness                     ; has one. Symptoms not explained
  creatinine blood increased       ; by these diseases are put aside.
Considering:                       ; These are the symptoms that are
  age 26 to 55                     ; explained by the diseases
Ruleout:                           ; that it is considering. It will
  hepatitis chronic                ; try to rule out all but one of
  alcoholic hepatitis              ; the diseases.
Abdomen pain generalized?          ; In order to do this it asks
* no                               ; for further information.
Abdomen pain right upper
      quadrant?
```

**Figure 1.3** Output from the Caduceus program