



**Calhoun: The NPS Institutional Archive**  
**DSpace Repository**

---

Faculty and Researchers

Faculty and Researchers' Publications

---

1988

# Artificial Intelligence through Prolog by Neil C. Rowe

Rowe, Neil C.

Prentice-Hall

---

<http://hdl.handle.net/10945/36984>

*Downloaded from NPS Archive: Calhoun*



Calhoun is a project of the Dudley Knox Library at NPS, furthering the precepts and goals of open government and government transparency. All information contained herein has been approved for release by the NPS Public Affairs Officer.

**Dudley Knox Library / Naval Postgraduate School**  
**411 Dyer Road / 1 University Circle**  
**Monterey, California USA 93943**

<http://www.nps.edu/library>

---

# Introduction

## What artificial intelligence is about

*Artificial intelligence* is the getting of computers to do things that seem to be intelligent. The hope is that more intelligent computers can be more helpful to us--better able to respond to our needs and wants, and more clever about satisfying them.

But "intelligence" is a vague word. So artificial intelligence is not a well-defined field. One thing it often means is advanced software engineering, sophisticated software techniques for hard problems that can't be solved in any easy way. Another thing it often means is nonnumeric ways of solving problems, since people can't handle numbers well. Nonnumeric ways are often "common sense" ways, not necessarily the best ones. So artificial-intelligence programs--like people--are usually not perfect, and even make mistakes.

Artificial intelligence includes:

- Getting computers to communicate with us in human languages like English, either by printing on a computer terminal, understanding things we type on a computer terminal, generating speech, or understanding our speech (*natural language*);
- Getting computers to remember complicated interrelated facts, and draw conclusions from them (*inference*);
- Getting computers to plan sequences of actions to accomplish goals (*planning*);
- Getting computers to offer us advice based on complicated rules for various situations (*expert systems*);
- Getting computers to look through cameras and see what's there (*vision*);
- Getting computers to move themselves and objects around in the real world (*robotics*).

We'll emphasize inference, planning, and expert systems in this book because they're the "hard core" of artificial intelligence; the other three subareas are getting quite specialized, though we'll mention them too from time to time. All six subareas are hard; significant progress in any will require years of research. But we've already had enough progress to get some useful programs. These programs have created much interest, and have stimulated recent growth of the field.

Success is hard to measure, though. Perhaps the key issue in artificial intelligence is *reductionism*, the degree to which a program fails to reflect the full complexity of human beings. Reductionism includes how often program behavior duplicates human behavior and how much it differs when it does differ. Reductionism is partly a moral issue because it requires moral judgments. Reductionism is also a social issue because it relates to automation.

## Understanding artificial intelligence

Artificial intelligence techniques and ideas seem to be harder to understand than most things in computer science, and we give you fair warning. For one thing, there are lots of details to worry about. Artificial intelligence shows best on complex problems for which general principles don't help much, though there are a few useful general principles that we'll explain in this book. This means many examples in this book are several pages long, unlike most of the examples in mathematics textbooks.

Complexity limits how much the programmer can understand about what is going on in an artificial-intelligence program. Often the programs are like simulations: the programmer sets conditions on the behavior of the program, but doesn't know what will happen once it starts. This means a different style of programming than with traditional higher-level languages like Fortran, Pascal, PL/1, and Ada | REFERENCE 11, .FS | REFERENCE 11 A trademark of the U.S. Department of Defense, Ada Joint Program Office. .FE where successive refinement of a specification can mean we know what the program is doing at every level of detail. But artificial-intelligence techniques, even when all their details are hard to follow, are often the only way to solve a difficult problem.

Artificial intelligence is also difficult to understand by its content, a funny mixture of the rigorous and the unrigorous. Certain topics are just questions of style (like much of Chapters 2, 6, and 12), while other topics have definite rights and wrongs (like much of Chapters 3, 5, and 11). Artificial intelligence researchers frequently argue about style, but publish more papers about the other topics. And when rigor is present, it's often different from that in the other sciences and engineering: it's not numeric but *logical*, in terms of truth and implication.

Clarke's Law says that all unexplained advanced technology is like magic. So artificial intelligence may lose its magic as you come to understand it. Don't be discouraged. Remember, genius is 5% inspiration and 95% perspiration according to the best figures, though estimates vary.

## Preview

This book is organized around the important central ideas of artificial intelligence rather than around application areas. We start out (Chapters 2-5) by explaining ways of storing and using knowledge inside computers: facts (Chapter 2), queries (Chapter 3), rules (Chapter 4), and numbers and lists (Chapter 5). We examine rule-based systems in Chapters 6-8, an extremely important subclass of artificial intelligence programs. We examine search techniques in Chapters 9-11, another important subclass. We address other important topics in Chapters 12-14: Chapter 12 on frame representations extends Chapter 2, Chapter 13 on long queries extends Chapter 3, and Chapter 14 on general logical reasoning extends Chapter 4. We conclude in Chapter 15 with a look at evaluation and debugging of artificial intelligence programs; that chapter is recommended for everyone, even those who haven't read all the other chapters. To help you, appendices A-C review material on logic, recursion, and data structures respectively. Appendix D summarizes the Prolog language subset we use in this book, Appendix E summarizes the Micro-Prolog dialect, Appendix F gives a short bibliography, and Appendix G provides answers to some of the exercises.

## Keywords:

*artificial intelligence*  
*natural language*  
*inference*  
*planning*

*expert systems*  
*vision*  
*robotics*  
*reductionism*

[Go to book index](#)