

The background of the cover is a dark blue field. Scattered across it are numerous red wireframe rectangles of various sizes and orientations. A series of thin white lines crisscross the field, some connecting the corners of the rectangles, creating a complex geometric pattern. The overall effect is one of abstract, logical structure.

# HOW MACHINES THINK

NIGEL FORD

A GENERAL INTRODUCTION TO  
ARTIFICIAL INTELLIGENCE ILLUSTRATED IN  
**P R O L O G**

TP18  
F711

8861095

附件8张 1至8已编号

# HOW MACHINES THINK

A GENERAL INTRODUCTION TO  
ARTIFICIAL INTELLIGENCE

ILLUSTRATED IN PROLOG

**Nigel Ford**

*University of Sheffield, UK*



E8861095

**JOHN WILEY & SONS**

Chichester · New York · Brisbane · Toronto · Singapore

Copyright © 1987 by John Wiley & Sons Ltd.

All rights reserved.

No part of this book may be reproduced by any means, or transmitted, or translated into a machine language without the written permission of the publisher

***Library of Congress Cataloging-in-Publication Data:***

Ford, Nigel

How machines think.

Bibliography: p.

Includes index.

1. Artificial intelligence. 2. Machine learning.

I. Title.

Q335.F67 1987 006.3 86-15953

ISBN 0 471 91139 9

***British Library Cataloguing in Publication Data:***

Ford, Nigel

How machines think : a general introduction to artificial intelligence.

1. Artificial intelligence

I. Title

006.3 Q335

ISBN 0 471 91139 9

Typeset by Photo-graphics

Printed and bound in Great Britain

8861095

# HOW MACHINES THINK



*This book is dedicated to  
my mother and father, and to  
Rosalind, Dominic, and Alasdair*

# CONTENTS

<b>Introduction</b>	1
<b>Part A: Basic processes</b>	5
1 Storing knowledge	7
2 New knowledge for old	15
3 Spreading diseases and tracing relations: recursion	21
4 Do it yourself	33
<b>Part B: Strategies for intelligent behaviour</b>	53
5 Putting your (mathematically) best foot forward	55
6 Plans and subplans	63
7 More sophisticated strategies	71
8 Probabilities	80
9 Frames of reference and expectations	90
10 A simple blackboard system	101
11 Learning	118
<b>Part C: Intelligent systems and their development</b>	131
12 Working systems	133
STRIPS: means—ends analysis	133
ABSTRIPS: planning at separate levels	136
MOLGEN: planning at interacting levels	139

SAVOIR: dealing with probabilities	141
HEARSAY: blackboard techniques	144
EXPERT-EASE: learning by induction	147
AM and EURISKO: automated discovery using frames	150
ACT: modelling the human mind	154
<b>13 Knowledge-engineering tools</b>	<b>158</b>
<b>Part D: Computers as intelligent life?</b>	<b>171</b>
POSTSCRIPT	185
Appendix 1 Programming in PROLOG: some 'working' illustrations	187
Appendix 2 'Probability' and 'lookout' programs	188
Appendix 3 Addresses of suppliers of software development tools	195
Bibliography	199
Index	204

# INTRODUCTION

- THINK!
- Could you ever consider a machine to be 'intelligent'?
  - If so, what would it have to be able to do to qualify?
  - If not, how would you define 'intelligence' to rule out machines?

Human beings have created many impressive tools which greatly surpass many of our natural abilities. With them we can travel great distances at great speed, lift huge masses, destroy on a vast scale. Yet because a tool can do certain things better than we can, we don't call it 'intelligent'. A car isn't intelligent because it can travel faster and further than we can run. A crane isn't intelligent because its ability to lift surpasses our own.

What, however, about tools which emulate and surpass our ability to perform *mental* rather than purely physical activities? Computers represent the most successful and widespread devices in this category. They can handle vast quantities of numerical and verbal symbols with impressive speed and accuracy, and few of us would disagree that they are extremely useful tools. Yet increasingly we can find statements from highly qualified individuals saying not only that machines can be intelligent but even alive! Some would even go as far as Geoff Simons, Chief Editor at the National Computing Centre (1983, p. 195), who considers that:

There is overwhelming evidence that we are now witnessing the birth of a new family of living species on Earth—and this must be seen as one of the momentous events in the history of life.

At both extreme and more moderate levels there are many voices raised in the argument for and against the acceptance of the notion of intelligence in



machines. This issue will be discussed in more detail in the final part of this book. For the moment what is really interesting is that the question of whether we can think of computers as intelligent, even alive, can and has been raised and sustained at a serious level. What this state of affairs means is that computers are beginning to emulate and surpass certain aspects of human *thought*; so much so that we must urgently turn our attention to a close examination of the nature, potential, even dangers of machine intelligence. In particular, those of us who are not computer specialists must nevertheless be able to achieve an informed understanding of what is being developed, and to add our voices to practical and philosophical aspects of the debate concerning how it can and should be applied. This book is aimed at the non-computer-specialist.

The range of mental activities which the computer can perform for us has greatly increased over recent years. Some of these are trivial or no more than routine in terms of what humans can do, although we have tended to associate them exclusively with human capabilities. Robots, for example, can learn a path through a maze or move objects around in seemingly intelligent ways. Other computer systems can converse with people, understand and write simple stories.

However, computers are also increasingly able to perform a range of what we think of as highly difficult and specialized intellectual tasks. Not only that—they are becoming as skilled and sometimes more skilled at them than their human counterparts.

So-called 'expert' computer systems already have an impressive list of achievements. We have examples of systems being used to great social and financial gain, as noted by Hayes-Roth *et al.* (1983, p. 6):

PROSPECTOR has discovered a molybdenum deposit whose ultimate value will probably exceed \$100 000 000 ...

R1 configures customer requests for VAX computer systems ... despite the fact that even the resident experts thought it could not be done ...

DENDRAL, which years ago demonstrated its superhuman performance, supports hundreds of international users daily in chemical structure elucidation ...

CADUCEUS embodies more knowledge of internal medicine than any human and can correctly diagnose complex test cases that stymie human experts ...

PUFF integrated knowledge of pulmonary function disease with a ... system for diagnostic consultations and now provides expert analyses at a California medical center.

High levels of performance have been achieved in other fields. Doug Lenat's AM (Automated Mathematician) system, for example, was designed to make mathematical discoveries. Soon after it started running it 'rediscovered' on its own many basic mathematical concepts and theorems, including Euclid's

Unique Factorization Theorem. It also discovered a number of interesting ideas unknown to its creator. Computer systems have now actually *created* new knowledge previously unattainable by human beings. Examples range from chess end-games to the diagnosis of heart-disease.

Interestingly, one of the most successful ways of devising 'intelligent' expert systems is not to design esoteric, highly complex computer procedures. Rather, it is to try to reproduce as accurately as possible in a computer system the relevant knowledge of human experts, gleaned from talking to and observing them. In this way, computers can not only 'think' as well as or better than humans, they can increasingly 'think' *like* humans.

For this reason, *human* thought will form both a starting- and a continual reference-point in this book in relation to the description and explanation of artificial intelligence—or the science of *machine* thought.

This book is designed for the non-computer specialist. Although details of programs are given (using the programming language PROLOG), it is possible to read these programs in plain English. English and PROLOG versions are given generally side by side. The main purpose of the book is not to teach you PROLOG (although you can, if you wish, pick up a very useful introduction to the language from it). Rather, it is to introduce many key ideas and principles used in work on artificial intelligence.

It tries to go some way towards solving a difficult problem in learning about artificial intelligence. You can often read about what sophisticated computer systems can do without any clear idea of *how* they do it. We hear of computers 'thinking', 'learning', using 'lines of reasoning', 'solving problems', and 'giving expert advice', but at a more basic level how can they do this? It's easy to see that computers can make calculations and print things on paper or on a screen—but in what way can they 'think'? To answer this, we must look at the details of programming. At this more detailed level it's often all too easy to get involved in learning about a programming language in great detail but be unable to see the wood for the trees. In other words, the sort of level at which you can understand programming is far removed from that at which sophisticated 'intelligent' computer systems operate.

There is bound to be a gap between the two levels, but this book aims to minimize that gap by building up an understanding of complex systems by first selectively looking at the details of programming. The examples of systems in Part C illustrate many of the concepts and ideas introduced at this more basic level in the earlier parts. This is made easier by the choice of programming language, PROLOG. This stands for PROgramming in LOGic, and PROLOG is much more closely related to how we think about problems than are the more usual languages such as BASIC, PASCAL, etc. In many important ways we can forget about the problems and intricacies of telling the computer exactly what to do at each stage of a program. Instead, we can concentrate much more directly on the nature of the task or problem the

program is meant to be dealing with. For this reason, the level at which the non-computer expert can understand the programming details, and the level at which complex, 'intelligent' computer systems work, are much closer together. For this reason, this book includes both levels. Hopefully, you will be able to see a fair amount of both the wood and the trees making up the wood!

PROLOG is the language which the Japanese have chosen to form the basis of their ambitious program to develop fifth-generation computers capable of highly intelligent behaviour. It is already an important artificial-intelligence programming language, particularly in Europe, and is likely to increase in importance. There is also growing interest in teaching PROLOG in schools.

There are a number of different versions of PROLOG available for both large and small computers, including a number of home micros. The version used in this book is the one which is easiest to read, a version developed at Edinburgh University. It is the nearest we have to a 'core' PROLOG— a universally accepted common form. Other versions differ mainly in the punctuation and layout, and, to a large extent, one can be easily translated to another.

This book includes parallel English versions of the programs, so if you are following the book in PROLOG as well as plain English, and have access to a different version of PROLOG, you can easily translate into your own version.

**Part A:**

# **BASIC PROCESSES**



## CHAPTER 1

# STORING KNOWLEDGE

A good starting-point in any search for comparisons and contrasts between human and machine thought is the question of *memory*. Man's collective memory—both in folklore and in libraries of recorded knowledge—is central to the development of his race. It allows each generation to function effectively on the basis of shared experience. It also allows each generation to rise increasingly higher on the shoulders of previous generations. In the same way, the memory of each person is vital to his or her ability to function and develop over time.

Memory is also of prime importance in the functioning and development of machine intelligence. Indeed, ways of representing knowledge in computer memory is a key area of research in artificial intelligence. Human beings can certainly represent large amounts of knowledge in their memories and use it flexibly and effectively. This is the basis for intelligent and creative thought and behaviour. Their weakness lies in their ability to forget and tire. Knowledge may often be partially or inaccurately remembered and applied. Efficient reasoning may not always be consistent.

Computers have complementary strengths and weaknesses. They can remember vast quantities of information and apply reasoning procedures accurately and tirelessly. However, their ability to use knowledge flexibly and creatively is (at present) more limited. Perhaps the true potential of artificial intelligence will lie not so much in the development of intelligent machines which will replace human expertise but rather in the linking of human and machine strengths and the minimizing of the weaknesses of each.

Let us first examine the nature of human memory. How do we store knowledge and how do we use it?

A common method employed by researchers to represent the contents of human memory is the *semantic network*. This consists of *nodes* (the dots in Figure 1) and *links* (the arrowed lines). A number of models of human memory have been developed which use semantic networks to represent knowledge. The close links between computer science and psychological research is




Figure 1. A semantic network

brought home by the fact that the first model of semantic memory—that of Ross Quillian—was constructed as part of a computer program, the ‘Teachable Language Comprehender’. This was designed to be able to understand and produce language. The psychological validity of Quillian’s model was subsequently tested and further models developed in the light of such studies. Two such models of memory are Anderson and Bower’s ‘Human Associative Memory’, and Lindsay, Norman, and Rumelhart’s ‘MEMOD’. The models are very different in many ways but all use a system of semantic networks. Examples of these networks are shown in Figure 2. *Don’t worry about the details.* The point is to illustrate different ways of using semantic networks made up of nodes and links. Much more complex ideas can also be represented in basically the same way, as illustrated in Figure 3.

- THINK!
- Translate the semantic network shown in Figure 4 into seven short sentences. Here is one: ‘John likes Anne’—you try the other six.
  - Now arrange the sentences in the following pattern:  
verb (‘subject’, ‘object’).  
For example:  
likes (‘John’, ‘Anne’).

Now turn the page upside-down to see my answer!

ENGLISH

Mary likes John  
John likes Anne  
Anne likes Fred  
John dislikes Mary  
Anne dislikes John  
Mary envies Anne  
John envies Fred

PROLOG

likes('Mary','John').  
likes('John','Anne').  
likes('Anne','Fred').  
dislikes('John','Mary').  
dislikes('Anne','John').  
envies('Mary','Anne').  
envies('John','Fred').

The semantic network you have just translated could be typed into a PROLOG system in the following form:



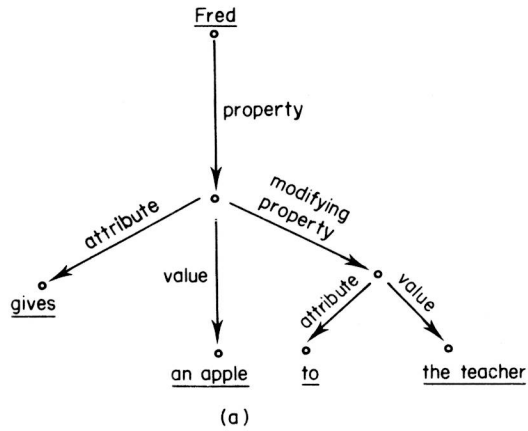


Figure 2. Models of how the sentence 'Fred gives an apple to the teacher' is stored in human memory. (a) Quillian's model; (b) Anderson and Bower's model; (c) Lindsay, Norman, and Rumelhart's model. (Adapted from Open University (1978). D303, *Cognitive Psychology*, Block 3, Units 18/19, p. 60)



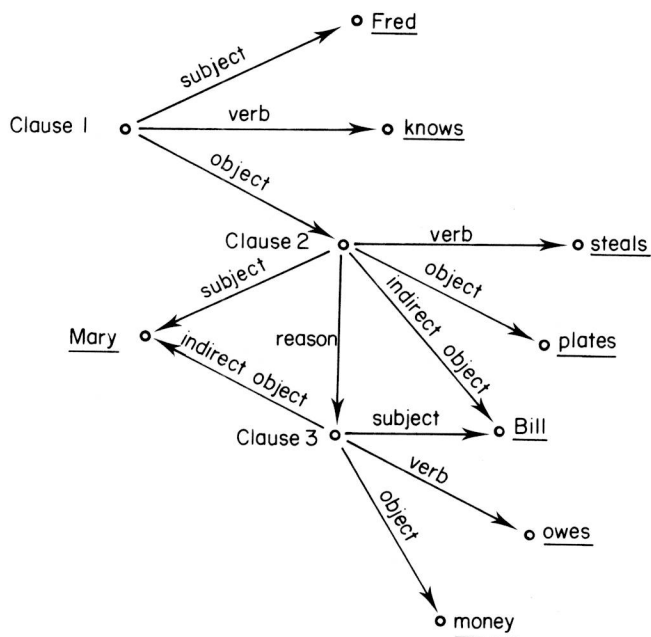


Figure 3. Semantic network representation of the sentence 'Fred knows that Mary steals plates from Bill because Bill owes Mary money'. (Adapted from Open University (1978). D303, *Cognitive Psychology*, Block 3, Units 18/19, p. 105)




Figure 4. Another semantic network