EXPERT SYSTEMS

AN INTRODUCTION FOR MANAGERS

ANNA HART

First published in 1988 by
Kogan Page Ltd,
120 Pentonville Road, London N1 9JN

British Library Cataloguing in Publication Data

Hart, Anna
Expert systems: an introduction for managers.
1. Expert systems
I. Title
006.3'3

ISBN 1-85091-368-4

Photoset in North Wales by
Derek Doyle & Associates, Mold, Clwyd.
Printed and bound in Great Britain by
Billing and Son Ltd, Worcester

EXPERT SYSTEMS

AN INTRODUCTION FOR MANAGERS

PREFACE

'The pen is a formidable weapon, but a man can kill himself with it a great deal more easily than he can other people.' George Dennison Prentice, Preticeana (1860).

'Oh no,' I hear the cry, 'not another book about expert systems', and certainly anyone who contemplates the plethora of books already published on the subject could be forgiven for thinking that the addition of yet another requires some justification. But, in my experience, the existing literature suffers from a number of shortcomings as far as potential users of expert systems are concerned. The majority of the books tend to be technical, describing how to build systems rather than describing in practical terms what is to be built.

Even those that are addressed to a wider audience often fail to address fundamental issues, and, in some cases are positively misleading and unrepresentative of what is really happening. Too often the authors have tacitly assumed that expert systems are obviously a good idea in almost any situation.

This book sets out to be non-technical and relatively easy to read, and above all, to discuss the practical issues that most concern the potential user, in a practical and straightforward manner. I hope that it achieves these objectives, but the reader is the final judge.

I also hope that it is realistic, honest and fair in its arguments. I hold and have often voiced strong and provocative views about expert system work. But in what follows I have tried to be unprejudiced and practical.

But I have also tried to tackle the technicalities head on. I have, of course, heard the argument 'Managers don't need to know all that. Just explain what an expert system is, and how much it will cost. That's all a manager needs to know in order to make a decision'. I do not believe this. For, as I try to show, it is far from easy even to provide an accurate description of 'an expert system' and it seems clear to me that successful management of innovation and change requires some understanding of the technology that underlies the change. Technical specialists are, however, notoriously bad at viewing technology in the wide context of organizations, and not much better at explaining technical matters to

the non-expert. It is, nonetheless, essential that managers should know which questions to ask and be able to judge the quality of the answers.

Having read this book managers, experts, computer users and computer specialists should have a good grasp of the terminology and issues associated with expert system development. Thus armed, they should be able to contribute to project development and evaluate progress. They will have a basic grasp of what is happening, and the role they are expected to play.

I was recently on a discussion panel at a conference introducing expert systems. The chairman likened expert systems to fruit on a tree: fruit can look ripe on a branch, but when it is picked it might be sour and useless. But if fruit is left too long it becomes over-ripe and no good. I ask the panel 'If apples are expert systems, is the fruit ready to pick and use?' Each of us gave a cautious reply in the affirmative. Then I added 'There are different types of apples. A ripe cooking apple will make a lovely pie, but is unsuitable for eating. Certain fruits are ripe but you need to know what they can be used for. I suggest that before picking you find out more about apples'. In this book I have attempted to provide some clues to those who wish to address that issue.

At another conference later that year system developers were complaining about the problems of educating management: the gist of their complaint was, 'We keep getting directives from people at the top to build systems, and it's clear that they don't understand about expert systems'. One speaker went on to state that a majority of companies were not actively involved in expert systems because of a lack of knowledge about what they offered. Again, I offer advice in these pages. The people who produce expert systems and the people who need them cannot just blame each other – we need to work together. If you want to find out more then read this.

The material comes from a variety of sources. First of all there is practical experience and knowledge of actual projects: tears, sweat, despair, hope and effort, though in some cases projects have been rendered unidentifiable for obvious reasons. Then there are the hours of discussion with people who have suffered my visits, or visited me, and numerous telephone conversations from people wanting advice. Some of these have been people who were coming to what I consider sensible conclusions, but seeking confirmation or confidence. I have also read much of the literature. It is not always possible to recall that an idea came from a particular book or paper, and failure to acknowledge such sources is an oversight, not deliberate. Finally, some of the material has already been used on courses we have run for local companies and, of course, undergraduate programmes.

I cannot claim any special qualifications, but I do find that convincing replies to some of my arguments are not always forthcoming. A student once described me as a Luddite, but I do not wish to be so. I do believe that people are important, and that computers exist for people and not vice versa. I am encouraged by those who thank me for trying to be honest.

To those readers who itch to corner someone like myself and ask for a simple explanation of 'the difference between the two types of chaining', I would say: please buy this instead and read Chapter 3. Whoever you are, do not accept anything (including what I say, but excluding this sentence!) without reasoning it out for yourself. Form your own opinions and be ready to justify them. Then please feel free to let me know your conclusions: that is how we all learn. One of the (few) advantages of being an author is receiving comments from happy or angry readers. I have included some approximate costs: these will need checking, as costs change so quickly.

Thanks are due to several people. My husband suffered the hours of bad temper and late meals, and even read an early draft. Such devotion is rare. Two students acted as a stimulus and encouragement: my thanks to Bruce Pilbrough and Peter Murray. I am indebted to the Computer Centre at Lancashire Polytechnic: had they not put the computer out of action over the summer I would not have had the time to write an original plan. A very stimulating week in Denmark gave further encouragement. In particular I thank Brian Mayoh for his vitality and originality, and Hanne and Steen Sondergaard for their fellow-spirit. Some sort of thanks are also due to a fearful virus which struck me down for months. This provided some time for writing, especially during convalescence. Ian Croall gave stimulating encouragement and although I cannot claim to understand half of what he knows I continue to learn from him. The final encouragement came from Lesley Ingleson, Chris King and Angela Murphy. Thanks for helpful criticism of the text are due to Lesley Ingleson and Phil Range.

Liz Giblin typed the text and converted wild scribblings into manuscript. Thank you. Kogan Page suffered my brainstorms. Piers Burnett made my English readable and did a marvellous job of editing. Martin Bilton delighted me with his perceptive comments and penetrating sarcasm.

May we all benefit. I may be culpable, but not for being silent.

Anna Hart

CONTENTS

1
SETTING THE SCENE

> The surest way to remain poor is to be an honest man.
>
> Napoleon I: *Maxims* (1804–15)

Although millions of pounds are currently being spent on the development of expert systems, and companies being told that they must invest in expert systems if they are to remain competitive, there are still certain basic questions which any potential purchaser of such a system will want to see answered before making what is likely to be a substantial investment:

- how do expert systems differ from other computer programs?
- what benefits do they offer?
- to what extent and in what fields have they already been proven?
- should our company invest in research and development?
- how are the systems built?
- what resources are needed to develop and then maintain the systems?
- how can a user manage the implementation of such a system in order to ensure satisfactory results?

This book tackles these questions.

1.1 Introduction

This book is about producing expert systems which are useful in that they will help companies to provide a more efficient or better quality service. It describes and examines some of the claims being made, raises questions, suggests some answers and offers explanations and advice. The intention is to give a realistic picture of what can be achieved, how, and why. The book does not teach programming, nor does it show how to design an expert system. There are very many such technical books on the market which cover these topics and another would add very

little to the existing pool of knowledge. What it does do is to consider practical problems based on case studies and actual experience of system development. Major interests are the production of high quality goods within reasonable budgets and timescales; the similarities and differences between conventional projects and expert system projects; the commitment which is needed, and the benefits which you can reasonably expect. The practical issues discussed in this book should be considered by any manager wishing to keep abreast of computer developments and available products. The material may also be of interest to other personnel involved in expert system projects; for example, potential users or experts.

1.2 The claims

I recently collected together a set of sales brochures which were being distributed by exhibitors at a conference for managers. The following is a fair synopsis of claims being made.

> 'There has recently been a revolution in the computer world, because artificial intelligence (A.I.) has moved out of the research laboratories into the marketplace. Any manager would be a fool not to take careful note of the products now on offer because they can increase efficiency and productivity with relative ease. These programs are not just better than conventional programs: they are totally different. They make the same decisions as a human expert, and they are accurate even when presented with inaccurate or incomplete information. They are able to weigh evidence and use judgement, and they do not make mistakes. They communicate in natural language so that professional expertise can be made available to almost anybody. The associated development tools can be used by a non-computer expert to build systems which can be used by a novice. Users love them. They even out-perform practitioners, and they have been proven in use over a number of years. You can save millions of pounds by investing in the products of this new revolution'.

Such claims sound irresistible and those who remain unconvinced by such hyperbole may well be swayed by the news media which have invested anything to do with the 'fifth generation' with an air of political urgency as well as an aura of technical glamour (Feigenbaum & McCorduck, 1984). Unfortunately, as we shall see later, research literature is sometimes misleading and there is, to say the least, some misunderstanding. Voices have already warned that certain products are being oversold or misused (Hart, 1986; Gotleib, 1985). But a critical examination of the claims would require some understanding of the

principles of expert systems as well as knowledge of products on offer or capable of being produced.

1.3 Why do I need to know about it?

The world is full of information: the problem is keeping track of those parts which are relevant to you. Computer science is particularly difficult, firstly because it is always changing, and secondly because it has its own impenetrable jargon. The jargon can be particularly confusing, and many books and magazines quickly degenerate into what is for many readers no more than technological mumbo-jumbo. Simple English words like 'field', 'block', 'track', 'bucket', 'word' and even 'bit' have their own meanings to computer scientists. Some words even have different meanings in different companies.

This terminology has recently been augmented by terms like 'intelligence, 'knowledge', 'reasoning' and 'judgement'. Such words are applied indiscriminately to both people and computers. This can be both suggestive and misleading. We all have some idea what is meant by a statement like 'he is obviously thinking about that' or 'you'll just have to use your own judgement'. But what does it mean to speak of a computer which 'thinks' or a program which exercises 'judgement'? In order to appreciate the tools and techniques which are associated with expert systems we have to fight our way through a morass of suggestions and allusions in order to gain the high ground where we may hope to find reasonable definitions and descriptions which adequately describe what is really possible.

The range of computer applications has changed over the past decade. The market place is flooded with products and salesmen. Computers and computer programs vary from the relatively simple and cheap and modest to the very sophisticated and astronomically expensive. Managers must keep in touch with this rapidly changing technology and monitor its progress so that they can assess the merits of new products and judge their relevance. Each new development needs to be assessed and its potential worth to the business evaluated.

Any tool can be useful. Unfortunately, in the past, many computers and programs have been bought and then misused, underused, or not used at all. Few people can afford to continue wasting money in this manner, especially as, in some ways, the more powerful the tools become the greater the danger of misuse. Anyone considering investment in expert systems should be able to fight through the maze of jargon, review the literature, learn to ask the right questions, assess potential application areas, and evaluate the available products.

Management should ensure that this is done, even if it is a delegated responsibility. If the decision is made to proceed then the buyer should be aware of what he can realistically expect as a return on his investment: if the decision is not to buy then he should know what is being refused and why. The task is not easy, but this text is designed to help. It cannot and should not prescribe a simple recipe for guaranteed success, but it can identify and discuss principles.

1.4 What is different?

It is perfectly true that there is a new range of software products available. The term 'expert system' is the name given to certain practical applications of artificial intelligence (A.I.). A.I. researchers define their aims as an effort to make computers perform tasks which would require intelligence if performed by human beings. The research is intended to:

- model intelligent behaviour,
- explain intelligent human behaviour,
- produce intelligent artefacts.

This does not help much. For what actually constitutes intelligence is very much a matter of opinion. In fact this issue was raised very early in the life of computers by the scientist Alan Turing, the inventor of the Turing test, and for the past 40-odd years, people have continued to speculate about how 'intelligent' computers could become, and whether they could properly be described as intelligent at all.

The answers to these questions are by no means obvious. For example, one could argue that the ability to perform high speed arithmetic is an indication of intelligence and this would imply that a pocket calculator was 'intelligent'. I think that this inability to define intelligence accurately has caused many of the misunderstandings associated with artificial intelligence. One A.I. enthusiast recently voiced his misgivings thus: 'as soon as we get a computer to meet a definition of intelligence the sceptics change the definition'. The point I am making is that while we have no difficulty in recognizing intelligent people by the things they can do we often feel far from happy when the same criteria are applied to machines. It is by no means clear that 'natural intelligence' and 'artificial intelligence' are comparable in any but the most superficial ways.

Certainly, bringing results from research laboratories into the market for industry and business has had varying degrees of success. Unrealistic aims have led to project failure, and misunderstandings

about what is possible or useful have caused confusion.

Let us return to the statements made in 1.2. Most of these claims are misleading if read in isolation from other material about computer systems.

To start with, the term revolution is inapplicable: evolution would be better. Certainly the use of new techniques has enabled computers to tackle different types of problems, but the programs are not necessarily easy to develop and, although they can be very beneficial and cost-effective, projects will only be successful if they are carefully planned and controlled. Moreover, as we seldom know exactly how human beings solve problems, we cannot say that programs perform in the same way. We may, in some situations, be able to produce programs that assist people in complex problem-solving, but inaccurate or incomplete information will lead to a possible answer – this is not surprising for there is no guarantee that a human being would be accurate in such a situation. Most expert systems *do* make mistakes.

Some of the development tools are relatively easy to use: none of the powerful ones is trivial to use. Building a powerful system for a novice is hard, and natural language facilities are almost always limited. Some systems have been proven in use; others have never been used commercially. While it is possible to save millions of pounds (see, for example, the discussion of the XCON project in Chapter 6), it is extremely easy to throw money away on an ill-conceived and poorly controlled project which will turn out to be worse than useless.

Expert systems are computer programs which are designed to manipulate information in a high-level way, and so emulate or assist human experts who employ expertise and knowledge. Knowledge is a very useful commodity. Programs which truly behaved like human experts would be very useful. They would not die, would not forget, would not lose their temper, etc.: in short they would have all the acknowledged strengths of computer programs without sharing many of the weaknesses of human beings. While some programs can do tasks like diagnosis better than some human beings in a well-defined area of problem-solving, others are far from being 'expert'. Even the 'expert' programs cannot outperform the human being in every respect.

We have come a long way in computer science, but there is still a long way to go before machines can begin to match many of the skills which we take for granted in our fellow human beings. Early computers did number crunching and were available to mathematicians who could code in low level computer languages. As languages like COBOL were developed computers moved into business, handling numbers and information commonly called data processing.

Now fast software and fast hardware enable computers to handle more complex information – this is called knowledge processing. There are new languages available, a range of packages (often called shells) are on sale, and you can build the systems yourself or get consultants. A new set of tools brings with it a new set of issues and questions. This book will not give all the answers, but it will discuss the important issues: you must make the final decision.