Ellis Horwood Series in **EXPERT SYSTEMS**

KNOWLEDGE ENGINEERING FOR EXPERT SYSTEMS

Mike Greenwell

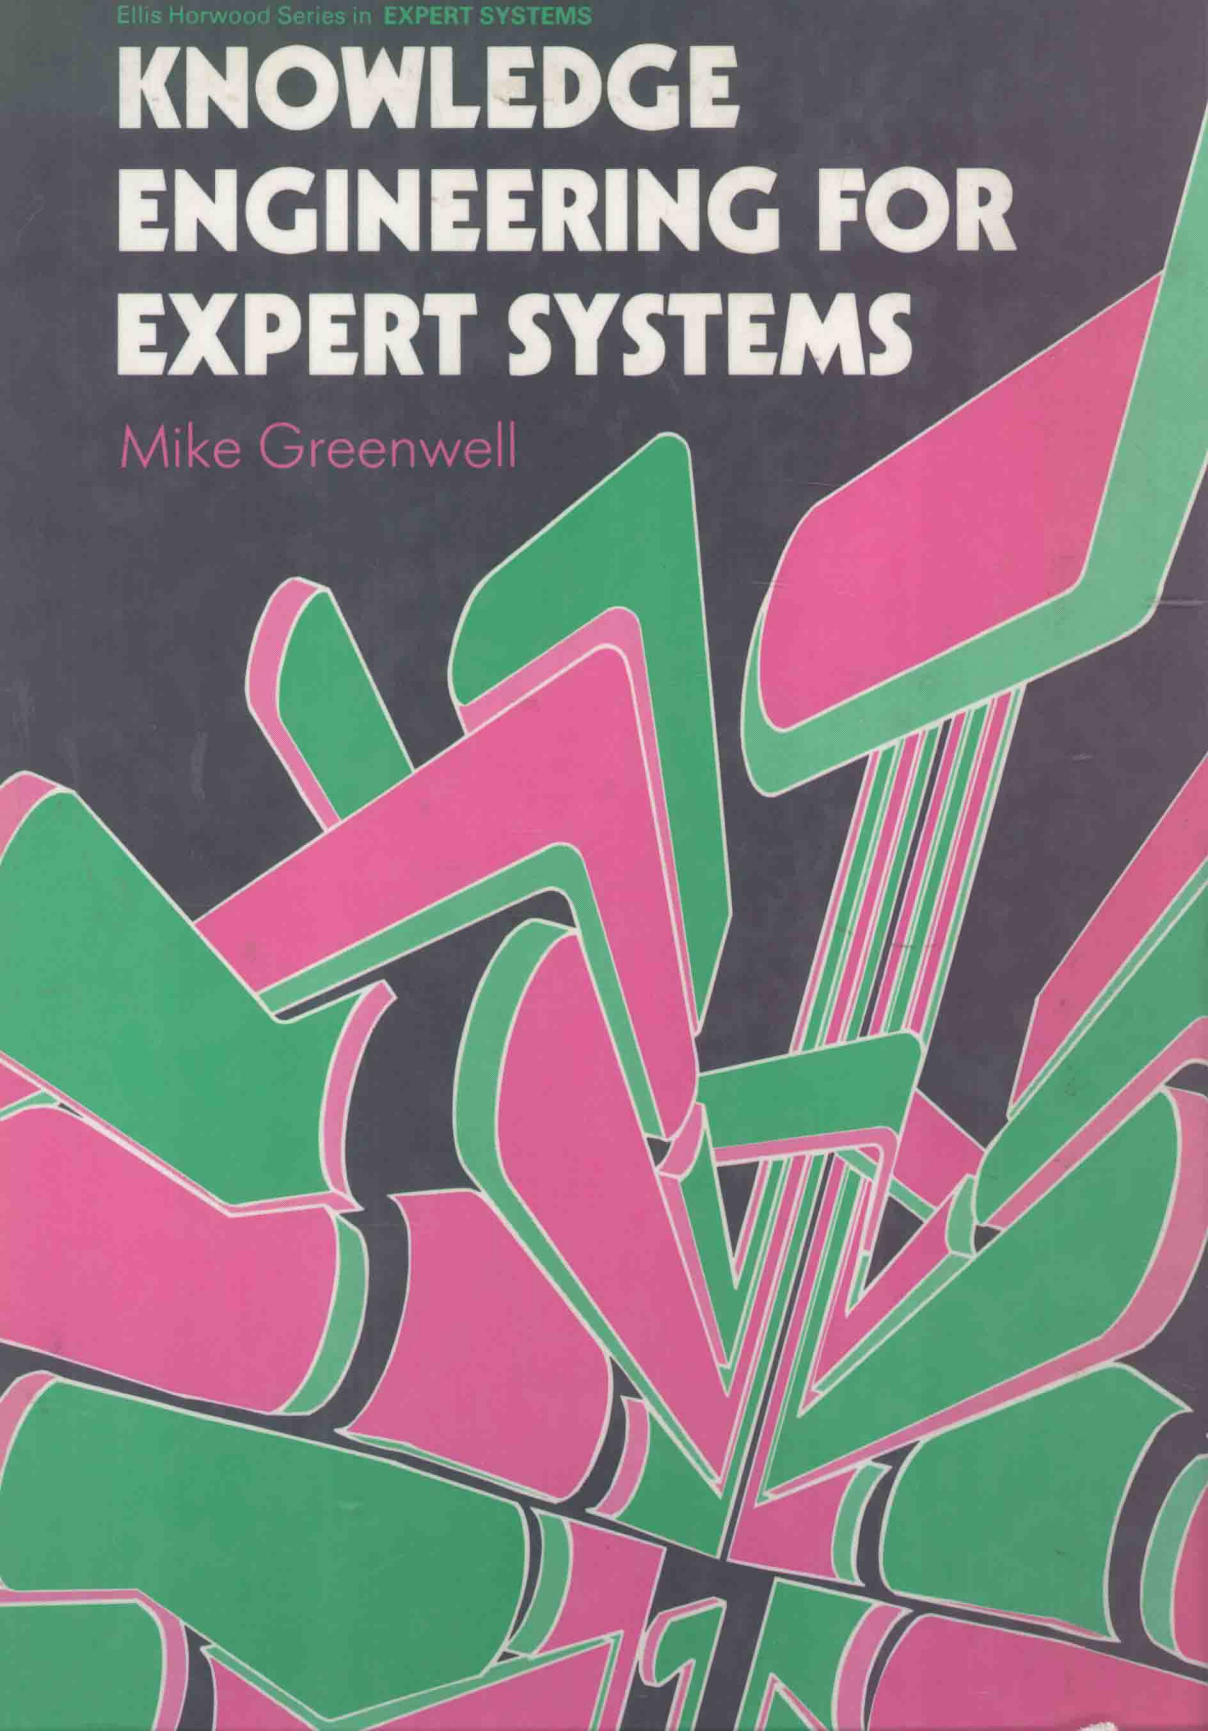

ELLIS HORWOOD BOOKS IN
INFORMATION TECHNOLOGY

General Editor: J. M. M. PINKERTON, Principal, J & H Pinkerton Associates (Consultants in Information Technology), and formerly Manager of Strategic Requirements, International Computers Limited

Ellis Horwood Series in
EXPERT SYSTEMS

KNOWLEDGE ENGINEERING FOR EXPERT SYSTEMS

MIKE GREENWELL, Independent Knowledge Engineering and Expert Systems Consultant

Expert systems have suffered from high costs, long development times, and have rarely produced the expected benefits. The major problems were the lack of clear techniques for the elicitation of knowledge and project management regime suited to the cyclic development of expert systems. This book shows how to reduce the costs attributable to the knowledge elicitation phase by offering a structured approach which includes a project management style that is both rigorous and flexible.

The book will help those knowledge engineers, expert systems project managers, and potential experts involved in a prospective project. Readers will be expected to know the meaning of such concepts as 'forward chaining rule' or 'expert systems tool kit', although a glossary has been included to refresh the memory. The author has worked directly in the field of expert systems since its rise to prominence, and the text reflects his expertise, providing a largely 'how to do it' approach. The basic material contained in this text is expected to remain valid for a long time, and benefits from real knowledge engineering experience (a rare commodity in this field).

The selection of topics is relevant and workmanlike, and the author's purpose is clear: how do you elicit knowledge from experts in a form suitable for expert system software, and then ensure that it has been done correctly and completely. He presents detail about knowledge engineering processes, outlining and linking it to later phases of development of expert systems, exposing the social and psychological factors affecting the knowledge engineer's job.

Readership: Knowledge engineering, information technology, expert systems, artificial intelligence, computer science, cognitive science.

KNOWLEDGE ENGINEERING FOR EXPERT SYSTEMS

ELLIS HORWOOD BOOKS IN INFORMATION TECHNOLOGY
General Editor: Dr JOHN M. M. PINKERTON, Principal, J. & H. Pinkerton Associates, Surrey (Consultants in Information Technology), and formerly Manager of Strategic Requirements, ICL

EXPERT SYSTEMS IN BUSINESS: A Practical Approach
M. BARRETT, Expertech Limited, Slough, and A. C. BEEREL, Lysia Limited, London

ELECTRONIC DATA PROCESSING, Vols. 1 and 2*
M. BECKER, R. HABERFELLNER and G. LIEBETRAU, Zurich, Switzerland

EXPERT SYSTEMS: Strategic Implications and Applications
A. C. BEEREL, Lysia Limited, London

SOFTWARE ENGINEERING ENVIRONMENTS
P. BRERETON, Department of Computer Science, University of Keele

SMART CARDS: Their Principles, Practice and Applications*
R. BRIGHT, Information Technology Strategies International Limited, Orpington, Kent

PRACTICAL MACHINE TRANSLATION*
D. CLARKE and U. MAGNUSSON-MURRAY, Department of Applied Computing and Mathematics, Cranfield Institute of Technology, Bedford

KNOWLEDGE-BASED SYSTEMS: Implications for Human–Computer Interfaces
D. CLEAL, PA Computers and Telecommunications, London, and N. HEATON, Central Computer and Telecommunications Agency, London

KNOWLEDGE-BASED MANAGEMENT SUPPORT SYSTEMS
G. I. DOUKIDIS, F. LAND and G. MILLER, Information Management Department, London Business School

KNOWLEDGE ENGINEERING FOR EXPERT SYSTEMS
M. GREENWELL, Expert Systems International, Oxford

KNOWLEDGE-BASED EXPERT SYSTEMS IN INDUSTRY
J. KRIZ, Head of AI Group, Brown Boveri Research Systems, Switzerland

ARTIFICIAL INTELLIGENCE: Current Applications*
A. MATTHEWS and J. RODDY, Aregon International Ltd., London

INFORMATION TECHNOLOGY: An Overview*
J. M. M. PINKERTON, J. & H. Pinkerton Associates, Esher, Surrey

EXPERT SYSTEMS IN THE ORGANIZATION: An Introduction for Decision-makers*
S. SAVORY, Nixdorf Computer AG, FRG

BUILDING EXPERT SYSTEMS: Cognitive Emulation
P. E. SLATTER, Telecomputing plc, Oxford

SPEECH AND LANGUAGE-BASED COMMUNICATION WITH MACHINES: Towards the Conversational Computer
J. A. WATERWORTH and M. TALBOT, Human Factors Division, British Telecom Research Laboratories, Ipswich

* In preparation

KNOWLEDGE ENGINEERING FOR EXPERT SYSTEMS

MIKE GREENWELL
Independent Knowledge Engineering
and Expert Systems Consultant

ELLIS HORWOOD LIMITED
Publishers · Chichester

Halsted Press: a division of
JOHN WILEY & SONS
New York · Chichester · Brisbane · Toronto

First published in 1988 by
ELLIS HORWOOD LIMITED
Market Cross House, Cooper Street,
Chichester, West Sussex, PO19 1EB, England
The publisher's colophon is reproduced from James Gillison's drawing of the ancient Market Cross, Chichester.

Distributors:
Australia and New Zealand:
JACARANDA WILEY LIMITED
GPO Box 859, Brisbane, Queensland 4001, Australia
Canada:
JOHN WILEY & SONS CANADA LIMITED
22 Worcester Road, Rexdale, Ontario, Canada
Europe and Africa:
JOHN WILEY & SONS LIMITED
Baffins Lane, Chichester, West Sussex, England
North and South America and the rest of the world:
Halsted Press: a division of
JOHN WILEY & SONS
605 Third Avenue, New York, NY 10158, USA
South-East Asia
JOHN WILEY & SONS (SEA) PTE LIMITED
37 Jalan Pemimpin # 05–04
Block B, Union Industrial Building, Singapore 2057
Indian Subcontinent
WILEY EASTERN LIMITED
4835/24 Ansari Road
Daryaganj, New Delhi 110002, India

British Library Cataloguing in Publication Data
Greenwell, Mike, *1953–*
Knowledge engineering for expert systems. —
(Ellis Horwood books in information technology).
1. Expert systems
I. Title
006.3'3

Library of Congress CIP available

ISBN 0–7458–0513–2 (Ellis Horwood Limited)
ISBN 0–470–21132–6 (Halsted Press)

Phototypeset in Times by Ellis Horwood Limited
Printed in Great Britain by Hartnolls, Bodmin

Table of contents

Acknowledgements

This is an eclectic work. I have drawn from my own experiences and from other authors in this area. As I have not always been able to find the original source of an idea I have been unable to acknowledge some of the authors directly. However, I do wish to acknowledge and give thanks to the following people whose work and ideas have influenced me and eventually led to this book: W. J. Clancey, M. D. Grover, whose early work was an inspiration to many, and Bob Wielinga and Joost Brueker of the University of Amsterdam. Their contribution to the study of knowledge acquisition has been my major influence since we were all involved on the same research contract. Many thanks to Margret Welbank for her encouragement and for her early review of the techniques. Many thanks also to those I have worked with who have influenced me, including Jonathan Killin of the KBSC, Ted Walker of ESI, from whom I learnt a great deal about building a real expert system, and finally Maureen Firlej, to whom I owe a special debt of gratitude for directing me towards knowledge acquisition back in 1984.

Many thanks to the following people for their assistance and advice during the seemingly never-ending task of writing this book: Chris Dobin, Masoud Yazdani, Alex Goodall, Louis Busuttil, Catherine Kenny, Paul Ormroyd, David Owen and the rest of the ALVEY EMEX club members for allowing me to use materials from the project.

1

Introduction

Expert systems are the limousines of the computer software industry. They are scarce and expensive. In recent years a number of researchers and knowledge engineers have been concerned to develop a methodology and thereby to reduce the costs involved in developing such systems. However, expert systems will always be expensive compared to conventional data processing, but it is hoped that this book and others now available will increase the reliability and utility of such systems as well as reduce the costs. By following the methodology presented here expert systems may be cheaper to build, which would mean they would be less scarce although still expensive compared to other information technology systems.

This book is not an introductory text to expert systems. There are many such books available if that is what is required. The target audience for this book are knowledge engineers, expert systems project managers and potential experts in a prospective project. The readers of this book are expected to know the meaning of such concepts as 'forward chaining rule' or 'expert systems tool kit'. However, a glossary has been included to refresh the memory.

Expert systems at the present time come in various shapes and sizes: the large application, the prototype investigation and the home-made application. The latter is produced by an individual with some expertise which he finds easy to represent using an inexpensive shell. The author is assured that this does take place but is unable to offer any evidence to support the claim, the probable reason being that successful home-grown expert systems are a prized and commercially sensitive resource for the individual in question, rather than that people are reluctant to admit failure in this enterprise.

The prototypic investigation is the most common variety of expert system. As a functioning piece of software it usually has the status of a toy and is used by management and technical staff as a vehicle to assist the general appreciation of expert systems technology. If the system develops beyond the prototypic investigation then it achieves the status of either a major commercial software system or a major project disaster.

A commercial expert system represents a substantial development cost for the large organisation which is prepared to bear the risk. It is for the technicians and management in those organisations that this book is intended. The technicians have become known as knowledge engineers, a title which seems to demand a greater respect than computer programmer or

systems analyst and has the advantage of being in common use even if it is somewhat nonsensical.

Expert system projects have a record of overshooting cost projections by an alarming margin. This is not a property of expert systems alone. Standard data processing projects have also had their failures. However, expert systems have the additional problem that, even when successful, the expected benefits are not always realised to the degree suggested by the advocates of the technology. A major problem in building expert systems is the lack of an integrated methodology for knowledge engineering and a complimentary managerial perspective. This book is one attempt to define a methodology for expert systems. The primary feature of that methodology is to control the most difficult and expensive phase of building expert systems — knowledge acquisition.

The task of obtaining knowledge, required to make expert systems function has been the least understood process in building knowledge-based systems. Indeed, some workers deny that knowledge acquisition is a major issue. This lack of interest in knowledge acquisition seems to stem from the conventional data processing background possessed by many of the workers in commercial expert systems. In recent years, however, more knowledge engineers have realised the difficulties inherent in acquiring knowledge from domain experts. Knowledge acquisition has become the major issue for knowledge engineers.

Academic criticism of early expert systems pointed to conceptual weaknesses in the knowledge-bases as a major problem. Solutions from researchers with backgrounds in psychology were largely concerned with psychological factors, which gave the task of knowledge acquisition prominence. While effort was being expended upon the knowledge acquisition/knowledge analysis stage, the commercial cutting edge of the technology was producing ever more elaborate software and hardware systems. The single user workstation and the AI tool kit or recommended programming environment became the principle tools of the knowledge engineer. Unfortunately, the principles for the acquisition and analysis of expertise lagged behind, truly a case of computer power without principles.

The expert systems industry is now entering a knowledge acquisition boom. Several books on knowledge acquisition have been published recently, training in knowledge acquisition is offered and specialist computer employment agencies are searching for candidates with knowledge acquisition skills for their clients. With all this activity it should be surprising to find that knowledge engineers are thrown into knowledge acquisition without preparation. And yet this still seems to be the case. One reason may be that knowledge acquisition looks easy. Indeed, what could be easier? — ask someone what they know, what they do, and how they do it, then plug all this into a ready-made expert system shell. In fact, one of the difficulties with knowledge acquisition is precisely this attitude. When problems arise, as they certainly will with this naive approach, the fault must be with the knowledge engineer, who seemingly is not competent at what is a very simple task.

Many of the issues concerned with knowledge acquisition fall into the domain of psychology. There has always been a close relationship between cognitive psychology and artificial intelligence. The cognitive psychologist takes an information processing perspective to mental phenomena. In artificial intelligence this is demonstrated by modelling human thought processes in a computer program. A cognitive model is not an expert system, although there are some parallels. Closer correspondence between the psychologist and the knowledge engineer is found in the identification of reasoning and problem solving strategies, memory constraints and the analysis of social situations such as the knowledge elicitation session itself. An awareness of psychology is an asset for the knowledge engineer but not as important as a detailed appreciation of programming styles relevant to the implementation of the knowledge analysis.

Knowledge engineering is that branch of the computer software industry concerned with building of expert systems. Knowledge engineering is itself composed of two separate tasks. On the one hand, there are knowledge engineers who mainly program in PROLOG or LISP.The software they build supports a particular knowledge representation, facilitating the inferencing mechanisms and providing a pathway to the user. An alternative would be to encode knowledge directly into the formalism offered by a sophisticated shell or using the facilities offered by a tool kit. Conversely, there are the knowledge engineers who spend much of their time talking to domain experts and potential users and analysing the transcripts of these conversations. Other tasks include designing the system and writing the code of which the knowledge-base consists. This second type of knowledge engineer is the primary audience for this book. That is not an invitation for the former to put this book down and concentrate on the latest work on object oriented programming. The greater the division in the labour of building expert systems, the greater the probability of communication problems which inevitably lead to a weaker final system.

Definitions should never be too rigid; the definitions with respect to knowledge engineering are developing as the principles and methodology develops. Some knowledge engineers may take issue with the definitions that appear below, Nevertheless, people require terms to be defined, however imperfectly.

Knowledge acquisition is the collection and analysis of information from one or more domain experts and any other sources leading to the production of a number of documents which form the basis of a functioning knowledge-base. A substantial subset of knowledge acquisition is knowledge elicitation. **Knowledge elicitation** is that area of knowledge acquisition which deals with acquiring information directly from domain experts.

Some writers on knowledge acquisition seem to be drawn into discussing the domain or a specific application, and confusing this with the knowledge engineering techniques. They make statements about the knowledge engineering which seem more properly to belong to the domain. I shall attempt to restrict this book to general knowledge engineering issues and keep references to specific domains to a minimum. However, examples from an

actual expert system project will be presented in chapter 11.

Instruction in the techniques of knowledge engineering will be found in chapter 5. Impatient readers are directed to move immediately to that chapter. Management issues are documented in chapters 9 and 10. Those with a specific interest in those issues are directed to that chapter and are also reminded that good management is founded on information, so they will be expected to also read those chapters dealing with techniques.

MODES OF EXPERT SYSTEMS

There are five major types of expert systems. The descriptions of each type are set out below for the purpose of providing background knowledge which is required for a number of planning decisions. This information should also form part of the introduction to expert systems for experts, which is given by the knowledge engineer as an ice-breaking exercise before the real work begins.

Knowledge-based information systems
Decision support systems
Consultation systems
Problem solving systems
Coaching systems

Fig. 1.1 — Basic types of knowledge-based systems.

Knowledge-based information systems tend to be intelligent-front-ends for accessing large databases. The knowledge required largely consists of the categorical structure of the items in the database and the user's typical modes of interaction with the database. The system assists with enquiry but does not solve problems or perform any recognised expert tasks. The knowledge required by this type of system is generally easy to obtain. Systems builders either have the requisite knowledge or can communicate without difficulty to those who have. Generally these are not very complex systems but probably beneficial and cost-effective.

Decision support systems take the advisory role of the expert. The system assists the user in the decision making processes by asking questions and interpreting the answers and has a powerful role as an organisational aid to decision making. However, there is little problem solving activity performed by the system. Even so, decision support systems tend to be fairly complex and require a lot of knowledge. The financial benefits may be low considering the probable high cost of building the system.

Consultation systems perform the reasoning in order to produce a diagnosis or report. A major feature is the ability for the user to interrogate the system's reasoning. However, explanation features are not very advanced as yet, and a large proportion of the work load will be concerned with the explanation facilities. Emphasizing explanations in the knowledge elicitation sessions will increase the potential power of the explanation

facility. Financial gain may not be realised in this category of expert systems if only because the failure rate of consultation system projects is so high.

Problem solving systems reason and offer a diagnosis with optional advice. Some problem solvers are not knowledge-based but fall within the category of artificial intelligence solutions. If the system is knowledge-based the knowledge engineer will need to acquire vast amounts of knowledge to reach some sort of performance parity with the expert. They are complex systems which run the risk of cracking under the weight of their own complexity. They can make or save a fortune for the organisation with the sophisticated hardware and software to develop and run the finished system.

Coaching systems are used to coach a student through the problem solving process. A feature of this type of expert system is the ability to adapt to individual users. The knowledge-base for a coaching system is probably not very large and may well consist of scripts which the user navigates through.

AUTOMATIC ELICITATION TECHNIQUES

In recent years there has been some work on the development of automatic systems for knowledge elicitation. This endeavour to replace the knowledge engineer by a computer program has been motivated by the difficulty and cost in eliciting knowledge. One of the reasons knowledge elicitation is so difficult and expensive is that approaches to it have been naive and unstructured. In such a situation programmatic simulations of knowledge elicitation have had to rely on procedures which were available to the programmer and possibly not used by the knowledge engineer. Three main thrusts seem to predominate: induction, reparatory grids, and documentation systems.

All three approaches have produced little in the way of useful tools for two main reasons. Firstly, the various knowledge representation languages and environments are usually incompatible with one another, so that a knowledge acquisition tool is usually dependent upon a particular environment which restricts the market for such a product. Secondly, the inferential power of these systems is particularly suspect and lacks the taken-for-granted common sense which is inherent in intelligent humans as they come to understand the expert's task and domain.

Knowledge engineering documentation systems seem to offer some useful organisational tools for the knowledge engineer. At their present level of development these systems are little more than intelligent word-processors with a powerful user interface and graphical capabilities. When a tool is useful it should be used; however, there are no recommendations in this book as to what packages the knowledge engineer should use. The remit here is limited to describing the knowledge elicitation situation and various techniques. In the near future there may well be powerful automatic elicitation systems. Until then, it appears that knowledge engineers will have to do most of the work themselves.

The rest of this book is organised in the following way:

Chapter 2: This is a highly speculative analysis of expertise, resulting in a basic principle which underpins the knowledge acquisition methodology outlined in this book. The purpose of this chapter is to stimulate if not enlighten the reader. There are many ways to explore the concept of expertise or knowledge; readers may like to develop their own analysis after reading this chapter.

Chapter 3: This is a comparison between knowledge elicitation and other forms of social interaction focussing upon the roles of the systems analyst and knowledge engineer.

Chapter 4: The subject matter for this chapter is interviewing theory relevant to knowledge elicitation. This is a subset of interviewing theory as a whole and therefore is a short but important chapter. The chapter concludes with an overview of the issues which concern the user of the expert system.

Chapter 5: This chapter describes in detail the major knowledge elicitation techniques.

Chapter 6: The first meeting and first interview between the knowledge engineer and the expert is very important. This chapter presents the content of the first meeting and a selection of questions which should be asked of the expert. Each question is discussed in terms of the reasoning underlying it and the ramifications for feasibility, design and management.

Chapter 7: This chapter deals with the planning and preparation issues involved in holding successful interviews, as well as the preparation and planning which is required for the project as a whole.

Chapter 8: The analysis of the knowledge elicitation interview and of the derived and verified information is the subject of this chapter. A number of intermediate representations are presented and discussed.

Chapter 9: The knowledge engineer writes the feasibility study using the information from the interviews outlined in chapter 6. This chapter describes the main issues which must be covered in that document.

Chapter 10: A recommended project management regime for expert systems is presented in this chapter. The reader is reminded that this area could quite easily form a book on its own. The orientation here is project management with emphasis on the knowledge acquisition.

Chapter 11: In order to pull together the ideas from the previous ten chapters a case study is presented. The case study in question is the ALVEY EMEX project.

2

The nature of expertise and knowledge acquisition

INTRODUCTION

The proposition that comprehending the nature of knowledge facilitates the building of expert systems is mistaken. Knowledge is a difficult concept to define. The constituents of knowledge, such as inference, experience, memory and certainty, require a complex conceptual analysis in their own right before the analysis of knowledge can proceed. The concept of expertise is more appropriate as it is more narrowly defined. This chapter describes a suggested framework for knowledge acquisition based upon an analysis of expertise. The veracity or utility of the principles underlying the framework are not important in their own right. Unlike the rest of this book the intention in this chapter is to explore some of the deeper issues. This is the only chapter which is oriented towards a more philosophical perspective. This chapter is here to be provocative and to be criticised. Its function is to make the reader think. The suggested knowledge acquisition framework is an incidental but nevertheless valuable result of the analysis. There is no requirement for knowledge engineers to consider at length what expertise is or any other related issues. However, knowledge engineers are by inclination curious and inquisitive people!

Expertise is directly related to what is done: it can be categorised in terms of an occupational category, such as lawer (legal), surgeon (medical) or social worker (caring services). Expertise has a psychological basis: it stems from memory, perception and problem solving skills and exists only in a social context which is, after all, where the expert system must reside. The analysis of expertise proceeds by asking, can an expert system be built to encapsulate any form of expertise? Are there some forms of expertise which would not be considered suitable for expert systems? This is really an enquiry into practicalities as well as the beginning of a philosophical discussion. The first step is to discover the mental and behavioural bases of expertise and then associate these with particular occupations.

FACTORING EXPERTISE

The analysis of expertise begins by identifying six factors, combinations of which may be used to describe any type of expertise which underlies a particular occupation. This does not imply that this is the only way of