


ELLIS HORWOOD BOOKS IN INFORMATION TECHNOLOGY


BUILDING EXPERT SYSTEMS

cognitive emulation

Philip E. Slatter



**BUILDING EXPERT SYSTEMS:
Cognitive Emulation**



ELLIS HORWOOD BOOKS IN INFORMATION TECHNOLOGY

General Editor: Dr. JOHN M. M. PINKERTON, Principal, McLean Pinkerton Associates, Surrey, (formerly Manager of Strategic Requirements, International Computers Limited)

STRATEGIC IMPLICATIONS AND APPLICATIONS OF EXPERT SYSTEMS

A. BEEREL, Director, Lysia Ltd., London

PRACTICAL MACHINE TRANSLATION

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

KNOWLEDGE-BASED EXPERT SYSTEMS IN INDUSTRY

J. KRIZ, Head of AI Group, Brown Boveri Research Center, Switzerland

ADVANCED INFORMATION TECHNOLOGY

J. M. M. PINKERTON, Principal, McLean Pinkerton, Esher

BUILDING EXPERT SYSTEMS: Cognitive Emulation

P. E. SLATTER, Product Designer, Telecomputing plc, Oxford

SPEECH AND LANGUAGE-BASED COMMUNICATION WITH MACHINES

J. A. WATERWORTH, British Telecom Research Laboratories, Ipswich

BUILDING EXPERT SYSTEMS: Cognitive Emulation

P. E. SLATTER, Ph.D.
Product Designer
Telecomputing plc, Oxford

Special Adviser:
Professor J. Campbell
Department of Computer Science
University College London



ELLIS HORWOOD LIMITED
Publishers · Chichester

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

First published in 1987 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

© 1987 P.E. Slatter/Ellis Horwood Limited

British Library Cataloguing in Publication Data

Slatter, P. E.

Building expert systems: cognitive emulation. —
(Ellis Horwood books in information technology).

1. Expert systems (Computer science)

I. Title

006.3'3 QA76.76.E95

Library of Congress Card No. 87-3760

ISBN 0-7458-0065-3 (Ellis Horwood Limited)

ISBN 0-470-20891-0 (Halsted Press)

Phototypeset in Times by Ellis Horwood Limited
Printed in Great Britain by R. J. Acford, Chichester

COPYRIGHT NOTICE

All Rights Reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without the permission of Ellis Horwood Limited, Market Cross House, Cooper Street, Chichester, West Sussex, England.

Contents

| | |
|--|-----|
| Preface | 7 |
| 1 Introduction | |
| 1.1 What is cognitive emulation? | 9 |
| 1.2 Current status of the emulation approach | 11 |
| 1.3 Objectives | 12 |
| 1.4 Scope | 13 |
| 1.5 Preview | 14 |
| 2 Human expert thinking | |
| 2.1 Introduction | 16 |
| 2.2 Cognitive psychology and the study of expertise | 18 |
| 2.3 Methods and techniques | 20 |
| 2.4 Research findings | 24 |
| 2.5 Summary and conclusions | 41 |
| 2.6 Suggested reading | 43 |
| 3 Arguments for and against emulation | |
| 3.1 Arguments for cognitive emulation | 44 |
| 3.2 Arguments against cognitive emulation | 48 |
| 3.3 Conclusions | 50 |
| 4 Factors facilitating and constraining cognitive emulation | |
| 4.1 Constraints on cognitive emulation | 52 |
| 4.2 Factors facilitating cognitive emulation | 56 |
| 4.3 Summary and conclusions | 61 |
| 5 Applications of the emulation principle: a survey of approaches | |
| 5.1 General introduction | 63 |
| 5.2 Emulating individual experts | 66 |
| 5.3 Emulating domain expertise | 72 |
| 5.4 Emulating expert cognition | 79 |
| 5.5 Emulating user cognition | 87 |
| 5.6 Emulating human information processing | 92 |
| 5.7 Emulating neural processing | 106 |

6 CONTENTS

| | |
|------------------------|-----|
| 5.8 General discussion | 109 |
| 5.9 Suggested reading | 115 |

6. Conclusions

| | |
|---------------------------------|-----|
| 6.1 Summary of main conclusions | 116 |
| 6.2 Summary of design advice | 119 |

| | |
|-----------------|-----|
| Glossary | 123 |
|-----------------|-----|

| | |
|-------------------|-----|
| References | 126 |
|-------------------|-----|

| | |
|---------------------|-----|
| Author index | 141 |
|---------------------|-----|

| | |
|----------------------|-----|
| Subject index | 144 |
|----------------------|-----|

Preface

This book is primarily intended for *expert system practitioners* considering whether to adopt cognitive emulation as a fundamental design principle. More generally, it is intended for designers wishing to tackle the emulation issues that arise in expert system projects in a more informed way.

Considering also the emulation of human thinking in knowledge-based expert systems, it is, additionally an exploration of the interface between knowledge engineering and cognitive psychology. While the differences between expert system design and psychological modelling have occasionally been discussed, a detailed assessment of the viability and implications of a cognitive approach to knowledge engineering has yet to be carried out. The present book addresses this issue, and is also intended as an introduction to the subject.

The following groups of people might also find the book of use:

Cognitive psychologists looking to evaluate the relevance of their subject to the expert system field, and to compare and contrast existing cognitive models of human expertise with applied expert systems.

Others: those involved in the introduction of expert systems technology in commercial and industrial settings should find at least Chapters 1, 3 and 4 helpful. Finally, persons interested in the debate over the comparative nature of human and artificial intelligence may find that this book throws a revealing sidelight on that controversy.

The treatment of psychological issues in this book assumes no prior exposure to cognitive psychology. And while a familiarity with the basic concepts and terminology of the expert system field is

assumed, a short Glossary of key terms is included to assist the more general reader. For introductory texts on expert systems Hayes-Roth *et al.* (1983) and Jackson (1986) can usefully be consulted. An extended introduction to modern cognitive psychology is provided in Anderson (1985).

This book has arisen from a research programme for a higher degree with the Open University. Over the last two years preliminary versions of parts of this book have appeared as a technical report, a conference paper, a review article (Slatter, 1985) and an Alvey Workshop paper (Slatter, 1986). Various individuals made valuable comments on these earlier efforts, and I would like to take this opportunity to collectively thank them. Some remarks by Ben du Boulay early on proved especially helpful. Thanks must also go to my supervisor at the Open University, Tim O'Shea. I am particularly indebted to John Fox, of the Imperial Cancer Research Fund, who acted as external supervisor in this research project. However, neither he, nor anyone else referred to, is responsible for the book's shortcomings — these are entirely attributable to the author. Finally, I would like to thank the chairman of Telecomputing plc, Bernard Panton, for making available the company's resources in producing this manuscript.

Philip Slatter

Oxford
January, 1987

1

Introduction

1.1 WHAT IS COGNITIVE EMULATION?

Expert systems use artificial intelligence (AI) techniques to solve problems ordinarily requiring human expertise. *Cognitive emulation* refers to a strategy in expert system design which seeks to emulate human thinking. A cognitive approach to knowledge engineering has several distinctive features:

- (1) It attempts to embody in an expert system not just the human knowledge of a domain expert, but also the way an expert represents, utilizes and acquires that knowledge.
- (2) The principle of cognitive emulation is usually defined to include the cognitive processes of system users, in addition to those of domain experts.
- (3) It is explicit. It enables the issues of expert and user emulation that may arise during an expert system development to be tackled in an explicit and principled fashion.
- (4) It attempts to emulate, using AI techniques, any aspect of human thinking that could assist in the construction of an expert system. This might include details of human memory organization, information processing limitations, problem solving and reasoning strategies, etc. It may also include emulation of the overall organization, or “architecture”, of human cognition.
- (5) It draws inspiration from empirical and theoretical investigations of human thinking — in particular, from research in Cognitive Psychology. The theories, hypotheses, computational models, methods and techniques of this branch of psychological science

are adapted for knowledge engineering purposes.

- (6) It is a concern with the practicalities of knowledge engineering (e.g. computational efficiency, modifiability, usability), which principally distinguish cognitive emulation from the cognitive modelling of psychologists.
- (7) At present, the influence of cognitive psychology on the expert systems field is essentially as depicted in Fig. 1.1, with the

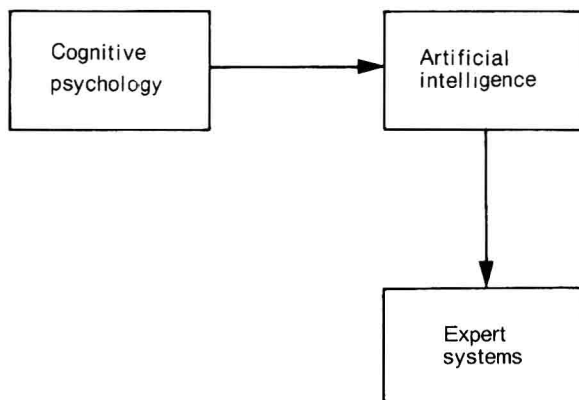


Fig. 1.1 — Existing indirect link between cognitive psychology and the expert systems field.

influence filtered through AI research. [AI scientists have capitalized on productive ideas derived from the study of human intelligence and developed them into a technology for creating artificial intelligence. Expert system designers employ this technology.] Adopting a cognitive approach implies supplementing this existing link with a more direct one, as shown in Fig. 1.2.

- (8) Finally, 'cognitive emulation' is both a descriptive concept and a prescriptive principle. As a descriptive concept, it can be argued that most expert systems incorporate — albeit unintentionally — many features characteristic of human knowledge processing (see Section 3.1.1). As a prescriptive principle, 'cognitive emulation' refers to expert system work in which an explicit strategy of emulating human cognitive processes is followed. Furthermore, the attempt should be based on some testable method, technique or model — that is, not solely on the casual observations or intuitions of the designer. It is with this latter, prescriptive definition of 'cognitive emulation' that this book is primarily concerned.

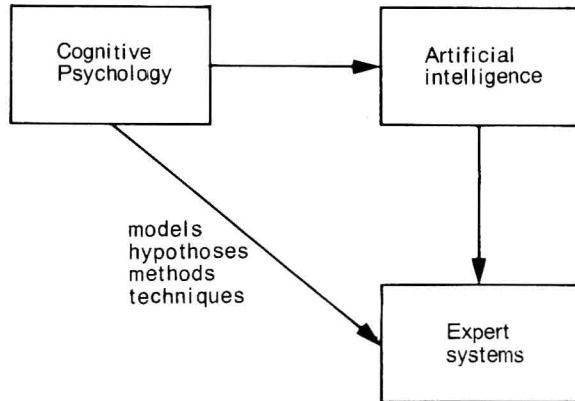


Fig. 1.2 — Additional direct link between cognitive psychology and the expert systems field implied by the cognitive emulation principle.

1.2 CURRENT STATUS OF THE EMULATION APPROACH

Throughout their brief history expert systems have been loosely modelled on the behaviour of human experts. However, systematic attempts at simulating experts' decision processes have only rarely been reported. Two early exceptions are: INTERNIST (Pople, 1982), which modelled the clinical reasoning of a diagnostician in internal medicine, and PSYCO (Fox, Barber and Bardhan, 1980), a system for diagnosing dyspepsia that incorporated various known principles of human information processing. The number of research projects involving a significant element of cognitive emulation has tended to increase in recent years (see, for example, papers in the volume edited by Coombs (1984) and Chapter 5 of this book). Similarly, the last few years have seen a growing awareness, among leading researchers, of the potential benefits of modelling human cognition more closely in expert systems (e.g. Clancey, 1984; Gaschnig *et al.*, 1983; Fox, 1982; Kuipers and Kassirer, 1984). Research is also underway to see how techniques developed in cognitive psychology can be used to facilitate knowledge acquisition from experts (e.g. Boose, 1984; Breuker and Wielinga, 1984; Gammack and Young, 1984).

However, despite this upsurge of interest, applications of the emulation principle in knowledge engineering are still the exception

rather than the rule. Most of the best known systems — including MYCIN (Shortliffe, 1976), PROSPECTOR (Duda, Gaschnig and Hart, 1979), DENDRAL (Buchanan and Feigenbaum, 1978) and R1/XCON (McDermott, 1982) — were constructed with little or no explicit aim of modelling expert thinking. Outside of research-oriented establishments there is, if anything, even less explicit concern with cognitive emulation. In commercial applications of today's expert systems technology the emphasis is firmly on achieving expert-level performance using formal problem solving methods. This viewpoint is expressed in the technology-based definition of expert systems offered by Johnson (1984, p.15) : “a set of computer programs which emulate human expertise by applying the techniques of *logical inference* to a knowledge base” (my italics).

One reason for the current status of the emulation strategy appears to be a general lack of appreciation among practitioners of the possibilities and issues associated with the emulation approach. But it could be argued that such an appreciation is not necessary. After all, if an explicit strategy of emulation is not a precondition of a successful expert system development — as systems such as DENDRAL and XCON appear to demonstrate — what justification can there be for adopting a cognitive approach? The pros and cons of emulation are dealt with in Chapters 3 and 4. For now, it is sufficient to point out that the objective of a well-articulated theory of knowledge engineering will only be achieved once all the relevant aspects of the subject — including cognitive emulation — have been properly evaluated. This book is intended as a contribution to this process of evaluation.

1.3 OBJECTIVES

This book is designed to perform a number of complementary roles:

- (1) *Introductory text* It is intended to introduce the subject of cognitive emulation in a way that is accessible to the wider knowledge engineering community — not just to research workers familiar with this field.
- (2) *Decision support* By making explicit the possibilities that exist for cognitive emulation, to help knowledge engineers make more informed decisions about adopting an explicit strategy of emulation in their own work. At the very least it should provide an informed basis for the handling of the emulation issues that can arise during an expert system development.
- (3) *Feasibility study* To investigate the viability of a strategy of

cognitive emulation by exploring in detail the issues that arise when a cognitive approach to expert system design is attempted. This requires consideration of such questions as:

- how do human experts actually solve problems?
- is cognitive emulation a theoretical possibility?
- what factors are likely to facilitate and constrain a cognitive approach?
- how should one decide when an emulation approach is worth adopting?
- what are the different approaches to emulation that have been tried?

- (4) *Information source* By supplying detailed references and an extensive bibliography of some 200 items, to provide the reader with the means to pursue particular lines of interest.

1.4 SCOPE

The scope of this book reflects these objectives:

First of all, the question of cognitive emulation is examined primarily from an (knowledge) “engineering” perspective. So, although the long-standing, and essentially philosophical, debate about the possibility of emulating human thinking in artificial systems is briefly referred to, this book is not directed to a further consideration of such meta-issues. Clearly, though, a book about cognitive emulation in expert systems is making a key assumption: namely, that in principle at least, computational models of human thinking are possible. Fortunately, this is an assumption shared by most workers in AI and cognitive science.

Second, this book is selective in which aspects of human cognition it deals with. Discussion focuses on cognitive processes relating to the “core” expert system topics of knowledge representation, knowledge acquisition, inferencing methods and system architecture. In the future, the human-like capabilities of expert systems are likely to be significantly enhanced by developments in such AI fields as natural language processing, vision and robotics. But these areas are not among the central concerns of knowledge engineers at present, and so are outside of the scope of this book.

The emulation principle is concerned with modelling human thinking in general, and the cognitive processes of selected categories of people in particular. Clearly, for the designers of computer systems intended to simulate expert-level performance, human experts are of primary interest. A whole chapter of this book is thus

devoted to reviewing our present understanding of human expert thinking — as revealed by research in cognitive psychology. The other group to be singled out consists of the users of expert systems. It is now widely acknowledged that the acceptance of an expert system can critically depend on the system being designed in accordance with the expectations, knowledge and preferences of its intended users. However, questions of user emulation can arise with most kinds of interactive computer system — not just expert systems. So it is not surprising that the subject of user cognition and its emulation have long been of interest to researchers in human-computer interaction (see, for example, Hammond and Barnard, 1985). The discussion here is restricted to user emulation as it relates to expert system design.

In general, the subject of cognitive emulation in this book is dealt with at a conceptual rather than an a detailed implementational level. So, in discussing applications of the emulation principle, attention is centred on the ideas and concepts involved, the success or otherwise of the project, and any general problems encountered in implementing cognitive constructs. For details such as the programming techniques employed, the reader should consult the supplied references. Similarly, the results of research in cognitive psychology are presented in a condensed format. Such a presentation is necessary here — but at the risk of masking the true complexity of the psychological issues involved. Again, the interested reader should consult the supplied references to obtain a fuller analysis.

1.5 PREVIEW

The remainder of this book is organized as follows :

Chapter 2 reviews psychological research on human expert thinking. This provides a baseline for comparisons with machine expertise in subsequent chapters. To provide some background for those new to this area, the study of expert thinking is placed in the context of modern cognitive psychology. Next, the principal methods and techniques employed by psychologists examining expert cognition are examined. The remainder of the chapter is given over to a review of the published literature on the nature and development of human expertise.

Chapter 3 reviews the main arguments for and against cognitive emulation in expert system design. The tentative conclusion reached is that a significant degree of emulation is inevitable, but that a pure, unselective strategy of emulation is neither realistic nor desirable.

Chapter 4 examines the prospects for cognitive emulation from a more pragmatic angle. Several factors are identified that represent constraints on the usefulness of a cognitive approach. Special emphasis is given to detailing areas of conflict with other knowledge engineering objectives. However, a second set of factors is identified which should facilitate an emulation strategy — especially in the longer term. Some guidance is given on when to seriously consider adopting an emulation strategy.

Chapter 5 is the longest chapter. It presents a critical survey of expert system research that has already addressed the emulation issue. Six basic approaches to cognitive emulation are distinguished and evaluated. This helps draw out in more detail the implications of an emulation strategy for knowledge acquisition, knowledge representation and system architecture. The chapter concludes by discussing the issues that arise when different approaches to emulation are combined. Some guidance is offered on how this might be achieved.

Chapter 6, the last chapter, summarizes the main themes and issues to have emerged. This is followed by a summary of the design advice contained in the book.

2

Human expert thinking

2.1 INTRODUCTION

This chapter presents an overview of research on expert thinking and how it develops. The principal source is research in cognitive psychology. The chapter has three main objectives:

- (1) *Appreciation*. It will indicate what a full commitment to a strategy of cognitive emulation in expert systems development could entail. This should provide an informed basis for the discussion of the emulation strategy in later chapters.
- (2) *Heuristic*. It attempts to fulfil a heuristic role — showing how cognitive psychology can function as a useful source of ideas in expert system design.
- (3) *Corrective*. It will aim to correct certain common misconceptions about the nature of expert thinking and what distinguishes it from novice thinking.

The following statements reflect some widely held beliefs about human expertise:

- Human expertise is acquired through experience.
- Human expertise is something mysterious and inexplicable.
- The superior performance of experts is based on superior intellectual ability.
- Experts reach conclusions by making a series of logical deductions based on the available evidence (Sherlock Holmes is a classic example from popular fiction of this view of the expert).