

Knowledge Acquisition: Principles and Guidelines

KAREN L. MCGRAW

KARAN HABBISON-BRIGGS

Knowledge Acquisition: Principles and Guidelines

**KAREN L. MCGRAW
COGNITIVE TECHNOLOGIES, ANNAPOLIS, MD**

**KARAN HARBISON-BRIGGS
UNIVERSITY OF TEXAS AT ARLINGTON**



PRENTICE HALL, ENGLEWOOD CLIFFS, NEW JERSEY 07632



© 1989 by Prentice-Hall, Inc.
A Division of Simon & Schuster
Englewood Cliffs, New Jersey 07632

All rights reserved. No part of this book may be reproduced, in any form or by any means, without permission in writing from the publisher.

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

ISBN 0-13-516436-2

Prentice-Hall International (UK) Limited, *London*
Prentice-Hall of Australia Pty. Limited, *Sydney*
Prentice-Hall Canada Inc., *Toronto*
Prentice-Hall Hispanoamericana, S.A., *Mexico*
Prentice-Hall of India Private Limited, *New Delhi*
Prentice-Hall of Japan, Inc., *Tokyo*
Simon & Schuster Asia Pte. Ltd., *Singapore*
Editora Prentice-Hall do Brasil, Ltda., *Rio de Janeiro*

Preface

It is with considerable pleasure that I accepted the offer to write the preface for Knowledge Acquisition: Principles and Guidelines. At the same time, the task holds considerable trepidation for me. For just as its authors have taken the risk of exposing their professional beliefs for the first time, so too, am I--this is my first preface.

I believe the professional knowledge engineer, as well as those aspiring to become such, will find the hefty work of Knowledge Acquisition: Principles and Guidelines to be very useful. After reviewing various versions of it a total of three times, I still find several gems tucked away. The first section contains guidelines which have "long" been part of the expert system development methodology folklore. How refreshing it is to see them collected, codified, and backed up by the authors' considerable professional experience, as well an unparalleled bibliography. Of special interest to me is the work detailed regarding the difficulties of garnering expertise from multiple experts, and some solutions.

The later chapter on tools is an excellent place to begin research, so long as the reader understands that this is an extremely rapidly changing arena. But the authors look beyond transitory features and trends and concentrate on the underlying issues and characteristics of a wide variety of tools available for the professional. I am particularly encouraged to see such attention paid to software tools. Just because we are "professionals" hardly means we are forced to rely solely on pencil, paper, and those "high-tech" tools of audio and video recorders! As a practicing software professional for 25+ years, it has always pained me how we consistently build very powerful and useful tools for others, yet persist in using paper and pencil when trying to do the critical work of needs assessment, requirements analysis, and specification. Only in the last few years has the concept of applying powerful software tools to the front-end of the software life cycle begun to make its appearance in the guise of Computer-Aided Software Engineering (CASE). This was possible only after some common methodologies were increasingly widespread (i.e., Yourdon/DeMarco and Gane/Sarson).

By the same token, I can see that the methodology espoused in Knowledge Acquisition: Principles and Guidelines may well begin to lay substantial groundwork for the field of knowledge engineering. Others will take up the challenge of building the automated tools to support the methods presented. So it is exciting to think that this work may have such an influence.

The initial primary audience for this work will probably be those in the software engineering world. Hence, these individuals will find the chapters on the subject to be a good introduction to the world of knowledge engineering. The analogies to software engineering are manifold, with considerable contributions from the field. By the same token, software engineers (and others) who take the time to read, understand, and discuss the material presented here will find their set of skills expanded considerably. More importantly, I believe, is that the reader will be able to formulate solutions for problem

domains once considered intractable. That, perhaps, is the essence of expert knowledge-based systems and the associated critical component of knowledge acquisition: that we will now be able to attack an increasingly complex set of problems which confront the world.

In summary, this book will be a key addition to the library of a professional. It is not intended to be an introductory text for the field of expert systems. There are many such excellent works. Nor is it simply a cataloguing of several papers on knowledge acquisition. It is a comprehensive, detailed, and first-hand account of an unusually literate professional on a very timely topic. I am sure the reader will agree!

Kenneth L. Modesitt
Chatsworth, California
May 20, 1988

Acknowledgments

Developing a book that reflects an integration of research and current development efforts requires the assistance of many professionals. Initially, we wish to thank the many knowledge engineers and domain experts who have worked with us on both small and large expert system development projects. We have taught each other much--and many of our shared "lessons learned" lie between the covers of this book.

In finalizing the contents of this book, we gratefully acknowledge assistance from colleagues in both university and corporate environments. In particular, we appreciate the suggestions and/or helpful reviews that the following individuals contributed:

Dr. Ken Modesitt, Rocketdyne, a Division of Rockwell International
Major Steve LeClair, Air Force Wright Aeronautical Laboratories
Dr. Robert Hoffman, Adelphi University
Dr. Nancy Martin, SoftPert Systems
Mr. John Boose, Boeing Computer Services
Dr. Bill Derr, Arizona State University
Ms. Kathy Palko, Texas Instruments
Mr. Stephen Kennedy, Grimes Insurance
Ms. Mary Seale, Texas Instruments and North Texas State University
Ms. Barbara Brown, Texas Instruments
Mr. Vince Waldron, Ohio State University and Waldron Communications
Ms. Kathy Waldron, Waldron Communications
Mr. Bruce McGraw, Falcon Microsystems
Mr. Artie Briggs, the H-B Group

Additionally, we salute the members of our families and the people with whom we currently work--both groups have displayed extraordinary patience and support, without which we could not have finished this project. Finally, we have benefitted from the editorial expertise of Ms. Valerie Ashton, Ms. Lisa Garboski, and Ms. Joyce Turner, our editors at Prentice-Hall, and Mr. E. David Zwieg! and Cognitive Technologies, who prepared the text for publication using Macintosh™ computers and Quark XPress™.

About This Book

The fascination with expert systems applications is evidenced by the appearance of expert system-related articles not only in academic journals but, more recently, in the business press. Some researchers (Partridge, 1986) claim that in the very near future, operational expert systems will be able to offer advice ranging from assistance in planning a complex business trip to sophisticated medical diagnosis.

Perhaps the prevailing enthusiastic view is based on the increasing number of successful expert systems. Expert systems are helping maintenance personnel diagnose specific types of system failures, offering financial planning advice, assisting with resource management problems, and even configuring complex computer hardware. In addition, the enhanced memory and speed of personal computers (PCs) and the concurrent development of PC-based expert system shells have propelled expert system technology out of the labs and into the marketplace.

However, problems in the development, maintenance, and enhancement of expert systems may severely restrict their integration into operational settings. While some of these problems relate to hardware deficiencies, such as the need for "real-time" processing, we believe that many others involve faulty knowledge-base development methodologies. In fact, the developmental area most often cited as the "bottleneck" in expert system development is knowledge acquisition, the process of extracting and translating expert-level knowledge into rules that become the heart of an expert system.

In this book we propose an organized approach to knowledge acquisition. This approach is based on experiences building both large and small expert systems in research and industrial settings. Within it we detail a set of techniques to help the knowledge engineer structure the knowledge acquisition process, integrating it into the expert system development paradigm. To increase the usefulness of this methodology, we offer guidelines, procedures, and a set of techniques for knowledge acquisition. Finally, we address the problem of evaluating the effectiveness of both the knowledge acquisition session or process itself and the resulting knowledge base.

Purpose and Content

This book presents a practical view of the knowledge acquisition process, its methodologies, and its techniques, in order to enable readers to develop expert system knowledge bases more effectively. It is intended to strike a balance between presenting (1) summaries of research in the field of knowledge acquisition and (2) methodologies and techniques that have been applied and tested on numerous programs in various contexts.

Intended Audience

This book is written for novice knowledge engineers or others tasked with acquiring knowledge for the systematic development of expert systems. Since knowledge engineers tend to come from diverse backgrounds, the presentation of the material does not presume a background in either computer science or artificial intelligence.

Project managers for expert system development projects may also be interested in this book, as it suggests a methodological approach to knowledge acquisition and presents chapters on guidelines and techniques that knowledge engineers can use to accomplish their tasks more effectively.

Structure

This text is structured in two parts. Part One includes Chapters 1 through 3 and presents background information to help the reader develop a successful knowledge acquisition methodology and program. Part Two, consisting of Chapters 4 through 11, describes suggested knowledge acquisition techniques, together with guidelines and considerations for their use.

Advance Organizers

Techniques to increase a reader's readiness for a chapter's content include the use of a Contents Page, Introduction, Objectives, and Key Terms.

Contents Page. Each chapter begins with a structure resembling a table of contents. Its purpose is to provide the reader at a glance with the contents and major focus of the chapter. Scanning the contents page provides the reader with a mindset for the material that will be presented. In addition, readers who are interested in specific sections may skim through the chapter to access them directly.

Introduction. The Introduction provides a context for the topics that will be covered in the chapter, a brief summary of the primary issues that will be discussed, and a general overview of the topics that will be presented.

Objectives. Each chapter was written with a specific set of purposes in mind. These purposes are enumerated in the Objectives section at the beginning of each chapter. Readers may wish to review the objectives prior to reading a chapter in order to select those chapters that best meet their needs. Reviewing the objectives upon completion of the chapter can help solidify the learning process.

Key Terms. The fields of expert systems and knowledge acquisition are laden with specialized terms. Terms that are critical to the comprehension of each chapter are presented at the beginning of each chapter. Each term is boldfaced when it is introduced in the text and it also appears in the Glossary at the end of the book.

Enrichment

The book is designed to encourage the reader to extend his or her understanding of the text through the presentation of Suggested Readings and opportunities for Application.

Suggested Readings. A brief bibliography of materials that provide more information on the topics presented in the chapter appears at the end of each chapter.

Opportunities for Application. Each chapter in Part Two (Chapters 4 through 11) concludes with an application activity. This activity provides an opportunity to extend the reader's understanding of the concepts presented in the chapter.

Chapter Contents

Chapter 1 presents background information on expert systems and the knowledge acquisition process. It concludes with a discussion of some major problems that affect knowledge acquisition to lay the groundwork for the philosophy and purpose of this text.

Chapter 2 reviews traditional software development methodologies and lessons that can be adapted for use in expert system development. It proposes a systems-oriented knowledge acquisition methodology that addresses some of the weaknesses that current literature ascribes to knowledge acquisition. This methodology provides the structure for the topics discussed in the remainder of the book.

Chapter 3 presents a practical, structured approach to knowledge acquisition, from using knowledge acquisition templates and building a knowledge acquisition database, to knowledge base review efforts. This approach includes tested guidelines and procedures that allow the knowledge acquisition process to be managed and controlled, enabling auditability and traceability of the developing knowledge base.

Chapter 4 acknowledges the crucial role played by the domain expert in the knowledge acquisition process and the importance of positive, effective knowledge engineer-domain expert working relationships. We suggest techniques for establishing and maintaining positive relationships, getting the most out of a knowledge acquisition session, and handling common problems.

Chapter 5 investigates the related problems of conceptualizing an expert's domain and analyzing the content of information garnered from knowledge acquisition sessions. It presents graphical, statistical, and other methods to portray domain information, either for use in structuring the knowledge base or as recording/feedback mechanisms.

Chapter 6 discusses the role of analysis techniques, such as task analysis, in a knowledge acquisition methodology. It describes a procedure for using analysis techniques to build a structure for the developing knowledge base and presents specific techniques knowledge engineers can use with different types of knowledge.

Chapter 7 examines the interview as a knowledge acquisition technique, explores how it can be made more efficient, and offers guidelines for planning and conducting knowledge acquisition interviews. It describes basic interview structures, various types of interview sequences, question types and levels, and techniques knowledge engineers can use in interviews to gather more usable knowledge.

Chapter 8 discusses process tracing and protocol analysis as techniques that can help knowledge engineers view and analyze a domain expert's decision-making strategies, priorities, and process. It provides some background on the use of these techniques in other fields, describe how to structure a knowledge acquisition session based on the use of these techniques, and explains how to use results from process tracing sessions.

Chapter 9 considers the rewards and problems associated with knowledge acquisition sessions with multiple experts. It presents guidelines for knowing when to use multiple experts, considerations for selecting multiple experts, techniques that are

appropriate for use with multiple experts, and suggestions for using knowledge acquired from multiple sources.

Chapter 10 focuses on the possibilities, appropriateness, and considerations for using automated knowledge acquisition tools, prototypes, and simulations for knowledge acquisition. It discusses the continuum along which knowledge acquisition may be automated and presents examples of some of the current crop of tools.

Chapter 11 explores an area that presents a major challenge to expert system effectiveness, the evaluation of knowledge acquisition effectiveness. It considers the verification-and-validation problem and presents some ongoing methods that developers can use to monitor the effectiveness of both the delivery of knowledge acquisition sessions and the content of the knowledge base.

An **Appendix** is provided for readers who desire more background information on theories related to how humans process and store information (Appendix A) and on disadvantages of verbal reporting (Appendix B).

The **Glossary** presents definitions for key terms used throughout the text.

Contents

PREFACE	ix
ACKNOWLEDGMENTS	xi
ABOUT THIS BOOK	xiii
1 KNOWLEDGE ACQUISITION FOR EXPERT SYSTEMS	1
Expert Systems: An Overview	2
Examining the Terminology	
Sampling Expert Systems	
Knowledge Acquisition in Expert System Development	8
Conceptualizing Knowledge Acquisition	
Knowledge Acquisition Modes and Approaches	
Knowledge Acquisition Process	
Identifying and Tapping Knowledge	13
The Nature of Knowledge	
The Relationship Between Knowledge and Expertise	
Impediments to Effective Knowledge Acquisition	18
Lack of Management and Organization	
Incompletely Trained Knowledge Engineers	
Translating Knowledge from Session to Code	
Knowledge Base Flexibility	
Handling Conflicting or Uncertain Information	
Verification and Validation Procedures	
2 A SYSTEMS ENGINEERING PERSPECTIVE ON THE KNOWLEDGE ACQUISITION PROCESS	29
The Process of Solving Conventional Problems	32
Software Engineering for Conventional Problem-Solving	
Knowledge Acquisition in Conventional Problem-Solving	
Knowledge Acquisition: Principles and Guidelines	iii

The Process of Solving Knowledge-Intensive Problems	40
Characteristics of Knowledge-Intensive Problems	
Software Engineering for Knowledge-Intensive Problems	
Documenting the Design Process	
Knowledge Acquisition for Knowledge-Intensive Problems	51
A Systems-Oriented Knowledge Acquisition Methodology	
An Adaptive Waterfall Model for Knowledge Acquisition	
 3 ORGANIZING A KNOWLEDGE ACQUISITION PROGRAM	 59
Getting Started	61
Domain Familiarization	
Knowledge Acquisition Facilities and Equipment	
Developing Knowledge Acquisition Procedures	70
Knowledge Acquisition Procedures	
Record-Keeping Procedures	
Review Procedures	
Orienting the Participants	86
Knowledge-Engineer Orientation and Training	
Domain-Expert Orientation and Training	
The Anatomy of a Session: A Case Study	89
Planning the Knowledge Acquisition Session	
Conducting the Knowledge Acquisition Session	
Reviewing the Session	
Session Follow-up	
 4 WORKING EFFECTIVELY WITH DOMAIN EXPERTS	 95
Building the Domain-Expert Pool	97
Factors Affecting the Selection Process	
Selecting the Initial Pool of Experts	
Selecting Specific Experts for Sessions	
Developing and Maintaining Rapport	106
Managing the Environment	
Managing the Session	
Common Problems	115
Analyzing Problem Behaviors	
Types of Problems	
Techniques for Difficult Situations	119
Reflective Listening	
Documenting Problems	

5	CONCEPTUALIZING THE DOMAIN	125
	The Role of Domain Conceptualization	127
	Initial Understanding	
	Estimating Knowledge Acquisition Needs	
	Knowledge-Base Design	
	The Challenge of Domain Conceptualization	
	Techniques for Evoking Domain Conceptualization	129
	Evoking Concept Identification	
	Techniques for Organizing and Analyzing Domain Conceptualizations	135
	Defining Concepts	
	Using Maps and Models	
	Identifying Relationships	
	Moving On	
6	STRUCTURING THE KNOWLEDGE ACQUISITION PROCESS	155
	Extending Systems Analysis to Expert System Development	157
	Knowledge Engineer as Systems Analyst	
	The Role of Analysis in Knowledge Acquisition	162
	Analysis Procedures for Knowledge-Base Development	
	Using Analysis Tools and Techniques	166
	Laying the Groundwork	
	Structuring the Domain	
	Completing the Analysis	178
	Refining the Initial Task Analysis	
	Planning Knowledge Acquisition Sessions	
7	INTERVIEWING FOR CONTENT AND CLARIFICATION	183
	The Interviewing Process	185
	A Model for Knowledge Acquisition Interviews	
	Interviewing for Knowledge Acquisition	
	Planning the Interview	
	Beginning the Interview	
	The Body of the Interview	
	Closing the Interview	
	Special Interview Skills and Techniques	196
	Using Questions Effectively	
	Levels of Questions	
	Sequencing Questions	
	Important Questioning Issues	
	After the Interview is Over	205
	Transcribing Session Information	
	Deciding What Is "Important"	
	Moving Towards Representation	

8 TRACING THE DECISION MAKING PROCESS TO ACQUIRE KNOWLEDGE	213
The Role of Decision Making in Knowledge Acquisition	215
Tracing the Decision Process	
The Usefulness of Verbal Reports	
Techniques to Compile Verbal Reports	
Managing Process Tracing in Knowledge Acquisition Sessions	221
Establishing a Methodology	
Sampling the Methods	
Extracting Knowledge from Process Tracing Sessions	230
Translating Sessions to Protocols	
Selecting the Analysis Framework	
Determining Decision Rules	
Refining the Analysis	
 9 ACQUIRING KNOWLEDGE FROM MULTIPLE EXPERTS	 241
Characteristics of Problems That Require Multiple Experts	243
Inconsistency Between Problem Domain and Available Expertise	
Subdivision of Knowledge Among Multiple Specialists	
Alternate Solution Sets	
Using Multiple Experts	245
Considerations for Using Multiple Experts	
Multiple Experts--Pros and Cons	
Approaches to the Use of Multiple Experts	251
Individual Consultation	
Group Consultation	
Techniques for Working With Multiple-Expert Teams	255
Brainstorming	
Consensus Decision Making	
Nominal Group Technique	
Computer-Facilitated Sessions	
Managing and Accommodating Multiple Knowledge Sources	264
Documentation	
Conflict Resolution	

10 AIDS AND TOOLS FOR KNOWLEDGE ACQUISITION	273
The Lure of Automation	275
Weaknesses of Manual Approaches	
Domain Expert as Knowledge Engineer	
Automating Knowledge Acquisition--A Continuum	281
Using Simulations for Knowledge Acquisition	
Using Prototypes of the System for Knowledge Acquisition	
Using Computer-Based Knowledge Acquisition Tools	
Types of Knowledge Acquisition Programs	284
Expert System Shells as Knowledge Acquisition Tools	
Tools Designed for Knowledge Acquisition	
Knowledge Acquisition Workbenches and Environments	
Problems to be Resolved	294
 11 EVALUATING KNOWLEDGE BASE DEVELOPMENT EFFORTS	 299
The Issue of Knowledge-Based System Evaluation	301
Verification and Validation	
Difficulties of V&V in Knowledge-Based Systems	
Attaining Verification of Knowledge-Based Systems	
Attaining Validation of Knowledge-Based Systems	
Evaluating Knowledge Acquisition Sessions	313
The Need for Ongoing Evaluation	
A Tool for Appraising Knowledge Acquisition Session Effectiveness	
 APPENDIX	 327
Appendix A: Human Knowledge, Memory, and Problem Solving	327
Appendix B: Disadvantages of Verbal Reporting	335
 GLOSSARY	 339
 REFERENCES	 351
 INDEX	 371

Knowledge Acquisition for Expert Systems

Expert Systems: An Overview

- Examining the Terminology**
- Sampling Expert Systems**

Knowledge Acquisition in Expert Systems Development

- Conceptualizing Knowledge Acquisition**
- Knowledge Acquisition Modes and Approaches**
- Knowledge Acquisition Process**
- Knowledge Acquisition Stages**

Identifying and Tapping Knowledge

- The Nature of Knowledge**
- The Relationship Between Knowledge and Expertise**

Impediments to Effective Knowledge Acquisition

- Lack of Management and Organization**
- Incompletely Trained Knowledge Engineers**
- Translating Knowledge from Session to Code**
- Knowledge Base Flexibility**
- Handling Conflicting or Uncertain Information**
- Verification and Validation Procedures**

INTRODUCTION

Knowledge acquisition is the most critical element in the development of expert systems. It is both people and time intensive. As expert system development finds its way out of research labs and into the business world, the cost of being people and time intensive will impact the success of a project. Current development methodologies stress the rapid prototyping approach, which provides few guidelines and procedures to help a knowledge engineer organize the knowledge acquisition process.

This chapter provides the reader with a foundation for the study of knowledge acquisition in expert systems development. This foundation helps one understand the "state of the art," typical problems, and critical issues. First, we provide an overview of expert systems, knowledge engineering, and knowledge acquisition, followed by a discussion of the various approaches or modes that have been used to acquire knowledge. Next, we present a historical view of knowledge acquisition stages. Subsequently, we provide information on the nature of knowledge, and we address problems associated with its identification and extraction. We discuss both the importance of knowledge and expertise in knowledge acquisition and the difficulties they present the knowledge engineer. In the final section we explore some of the impediments to expert systems development with which knowledge acquisition is associated.

OBJECTIVES

This chapter presents information on vocabulary, concepts, and problems related to knowledge acquisition. Specifically, it enables readers to accomplish the following:

- Understand basic terminology and concepts related to expert system development and knowledge acquisition
- Identify historical and current approaches to expert system development and knowledge acquisition
- Conceptualize the various categories of knowledge that must be acquired
- Identify difficulties in the acquisition of knowledge from domain experts
- Recognize impediments to the successful application of current approaches