

EXPERIENCE, MEMORY, AND REASONING

Edited by

Janet L. Kolodner
Christopher K. Riesbeck

7P11
421
8762683

EXPERIENCE, MEMORY, AND REASONING

Edited by

Janet L. Kolodner

Georgia Institute of Technology

Christopher K. Riesbeck

Yale University



E8762683



LAWRENCE ERLBAUM ASSOCIATES, PUBLISHERS

1986 Hillsdale, New Jersey

London

To Mike, Orly, and Joshua
Maxine, Michael, Elizabeth, and Jonathan

Copyright © 1986 by Lawrence Erlbaum Associates, Inc.

All rights reserved. No part of this book may be reproduced in any form, by photostat, microfilm, retrieval system, or any other means, without the prior written permission of the publisher.

Lawrence Erlbaum Associates, Inc., Publishers
365 Broadway
Hillsdale, New Jersey 07642

Library of Congress Cataloging-in-Publication Data

Kolodner, Janet L.

Experience, memory, and reasoning.

(The Artificial intelligence series)

Includes index.

1. Reasoning (Psychology) 2. Memory. 3. Experience
— Psychological aspects. 4. Artificial intelligence.
I. Riesbeck, Christopher K. II. Title. III. Series:
Artificial intelligence series (Hillsdale, N.J.)
BF441.K57 1986 153 85-20721
ISBN 0-89859-644-0

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

EXPERIENCE,
MEMORY,
AND REASONING

A Series of Monographs, Treatises, and Texts
Edited by Roger C. Schank

SCHANK AND ABELSON • *Scripts Plans Goals and Understanding:
An Inquiry into Human Knowledge Structures, 1977*

LEHNERT • *The Process of Question Answering:
Computer Simulation of Cognition, 1978*

MEEHAN • *The New UCI LISP Manual, 1979*

CHARNIAK, RIESBECK, AND McDERMOTT • *Artificial
Intelligence Programming, 1980*

SCHANK AND RIESBECK • *Inside Computer Understanding:
Five Programs Plus Miniatures, 1981*

KOLODNER • *Retrieval and Organizational Strategies
in Conceptual Memory: A Computer Model, 1984*

KOLODNER AND RIESBECK • *Experience, Memory,
and Reasoning, 1986*

List of Contributors

WILLIAM M. BAIN • Department of Computer Science, Yale University

B. CHANDRASEKARAN • Department of Computer and Information Science, Ohio State University

KURT P. EISELT • Department of Computer Science, University of California, Irvine

RICHARD H. GRANGER • Department of Computer Science, University of California, Irvine

JENNIFER K. HOLBROOK • Department of Computer Science, University of California, Irvine

JANET L. KOLODNER • School of Information and Computer Science, Georgia Institute of Technology

MICHAEL LEBOWITZ • Department of Computer Science, Columbia University

STEVEN L. LYTINEN • Department of Computer Science, Yale University

CHARLES E. MARTIN • Department of Computer Science, Yale University

DALE M. McNULTY • Department of Computer Science, University of California, Irvine

BRIAN J. REISER • Department of Psychology, Princeton University

- CHRISTOPHER K. RIESBECK** • Department of Computer Science, Yale University
- EDWINA L. RISSLAND** • Department of Computer and Information Science, University of Massachusetts
- ROGER C. SCHANK** • Department of Computer Science, Yale University
- V. SEMBUGAMOORTHY** • Department of Computer and Information Science, Ohio State University
- JEFF SHRAGER** • Department of Psychology, Carnegie-Mellon University
- ROBERT L. SIMPSON, Jr.** • Program Manager, Intelligent Systems, Defense Advanced Research Projects Agency
- ROBERT WILENSKY** • Computer Science Division, University of California, Berkeley

Preface

Several years ago, during a reunion banquet at the national Artificial Intelligence conference in Washington, DC, some of us were commiserating over the fact that the national and international AI conference had grown too big. The largeness and diversity of the conferences made it hard to do two things: to learn in detail about what others were doing and to present our own work at a satisfying level of detail. It seemed obvious that the thing to do was to begin having a series of smaller, more intimate thematic workshops devoted to presenting research ideas in detail.

The First Annual Workshop on Theoretical Issues in Conceptual Information Processing (TICIP) grew out of that. It was held in Atlanta, Georgia in March 1984 and included 50 people with roughly the same world view. In particular, all of us were interested in content-based theories of conceptual information processing. Despite this, our call for papers for the workshop elicited papers on a wide variety of topics: learning, parsing, creativity, problem solving, representation, and memory organization are just some of the topics that were discussed. Although topics were diverse, everyone was interested in a set of basic issues: What do people know? How do they use that knowledge to reason?

What we didn't know before we set up the workshop was that a majority of the papers would converge on a theme: the relationships between memory, experience, and reasoning. It has long been acknowledged that reasoning depends on heuristics or knowledge gleaned through experience. Thus, one relationship between reasoning and experience is already well-defined. We now add memory to the duo of experience and reasoning. How does memory fit in? The knowledge coming from experience that is used for reasoning must

be gleaned from an accumulation of many experiences. Those experiences do not necessarily come in the right order for learning. Yet, we, as people, learn from experience anyway. A memory stores experiences to allow learning to happen on the basis of several experiences. Furthermore, reasoning capabilities evolve as new knowledge is learned.

In addition, heuristic knowledge compiled from many experiences is not the only way we use experience in reasoning. Often, a *particular* previous experience acts as a guide in allowing us to construct a solution to a new problem. If we are *reminded* of a previous case similar to one we are trying to solve, we can use it as a guideline for solving the new one. Memory also stores experience, then, to be recalled to guide later processing. When we consider the role of a memory holding experiences, we also begin considering additional reasoning processes, namely those that consider previous individual cases to make new decisions. The name that has been coined for this view of intelligence is *memory-based reasoning*.

The chapters in this collection are based on presentations made at the First Annual TICIP Workshop. Each chapter addresses some issue associated with the relationships between memory, experience and reasoning. In considering representation, for example, the representation of episodes (experiences) is a major issue, as is their organization in a long-term memory for experience. Chapters in the understanding section consider the relationships between understanding processes and the memory store, including how we can expect knowledge to be organized in that store for access by understanding processes, how it can be updated as part of the understanding process, and what new processing considerations become important when we look at understanding in its relationship to memory. In the hypothesis generation section, a framework explicitly citing the relationships between problem solving, learning, and memory processes is presented, and other chapters explore issues associated with explaining why things have happened and justifying decisions, two reasoning processes that depend heavily on experience.

This book presents a first view of the issues associated with memory-based reasoning. There is much research currently being done along the lines presented in this book, both by those whose chapters are in this book and by others. This book presents an introduction to what we believe will prove to be an important approach to unraveling the mysteries of cognition and AI.

ACKNOWLEDGMENTS

Special thanks go to several people and agencies who helped with the First Annual TICIP Workshop: Shoshana Hardt, who was on the program committee along with us; Dana Eckart, Kirt Pulaski, Robert Simpson, and Katia Sycara, students at Georgia Tech who served as the local arrangements com-

mittee; the Georgia Tech School of Information and Computer Science, which provided their copying machines and administrative support; the National Science Foundation, who paid for secretarial time under Grant No. IST-8116892; and especially, Jean Early, who put in much time making sure everything ran smoothly. Funding that supported work on this book came from the National Science Foundation under Grant No. IST-831171, the Army Research Office under Contract No. DAAG-29-85-K-0023, and the Air Force Office of Scientific Research under Contract No. F49620-82-K-0010.

Janet L. Kolodner
Christopher K. Riesbeck

8762683

Contents



Preface xi

Chapter 1 Introduction 1

Knowledge and AI	1
Knowledge and Cognitive Psychology	2
Knowledge versus Memory	3
Problem Solving in Memory	7
Basic Questions	7
Organization of the Book	9

PART I: KNOWLEDGE ORGANIZATION AND REPRESENTATION

Chapter 2 Knowledge Representation—A Critique and
a Proposal 15
Robert Wilensky

Abstract	15
Introduction	15
The Problem With Conceptual Dependency	16
The Problem With Frames	18
KODIAK	19

	Primitive Epistemological Relations	21
	Examples	25
	Advantages of the Proposal	26
	Problems and Comments	27
Chapter 3	Learning and Memory in Machines and Animals: An AI Model that Accounts for Some Neurobiological Data	
	<i>Richard H. Granger and Dale M. McNulty</i>	29
	Abstract	29
	Introduction to the Problem	29
	Introduction to the Model: The 12 CEL Operators and Their Functions	32
	A Brief Example of the Operation of the CEL-0 Program	33
	Some Insights Resulting From Experience With CEL-0	38
	The Neurobiology of Multiple Memory Systems	41
	Conclusions: Artificial and Natural Learning Mechanisms	43
Chapter 4	Functional Representation of Devices and Compilation of Diagnostic Problem-Solving Systems	
	<i>V. Sembugamoorthy and B. Chandrasekaran</i>	47
	Abstract	47
	Motivation	47
	Components of a Functional Representation	49
	A Representational Scheme for the Functioning of Devices	49
	Compilation of a Diagnostic Problem-Solving System	56
	The Diagnostic Task and the Structure Produced by the Compiler	59
	"What Will Happen If" Problem Solving Using the Functional Representation	63
	Relation to Other Work	65
	Concluding Remarks	68
	Appendix	69

Chapter 5 Knowledge-Directed Retrieval of Autobiographical Memories 75
Brian J. Reiser

Abstract	75
Accessing Autobiographical Experiences	75
A Model of Directed Retrieval	76
Arguments for Directed Search	
Models	80
Psychological Evidence for Directed Search	86
Effective Search Strategies	90

**PART II: MEMORY-BASED HYPOTHESIS
FORMATION**

Chapter 6 Problem Solving and Dynamic Memory 99
Janet L. Kolodner and Robert L. Simpson, Jr.

Abstract	99
Experience's Roles in Problem Solving	99
Encoding and Organizing Experience in a Dynamic Memory	102
Experience Contributes to Change in Memory's Structures	104
Experience Contributes Exemplars for Analogy	107
Programs	112
Where Do We Go From Here?	113

Chapter 7 Learning How to Argue: Using Hypotheticals 115
Edwina L. Rissland

Abstract	115
Introduction	115
Uses of Hypotheticals	116
General Roles Played by Examples	122
Strategies of Using Hypos	123
Summary	125

Chapter 8	Assignment of Responsibility in Ethical Judgments	
	<i>William M. Bain</i>	127
	Abstract	127
	Introduction	127
	Background: The Task of the Judge	128
	A Mock Plea Bargaining Session	129
	Conclusions	137
Chapter 9	Explanation: A First Pass	
	<i>Roger C. Schank</i>	139
	Abstract	139
	Introduction	139
	The Role of Reminding	140
	Making Explanations	141
	The Explanation Process	143
	Reminding as Verification	143
	Finding Anomalies	144
	What Must Be Explained	147
	Types of Explanations	158
	Knowing When an Explanation is Right	161
	Conclusions	164
Chapter 10	Acquisition of Device Models in Instructionless Learning	
	<i>Jeff Shrager</i>	167
	Abstract	167
	Instructionless Learning	167
	What is Learned in Instructionless Learning?	169
	Hypotheses and Experimental Failure and Explanation	171
	The Structure of BigTrak Knowledge	173
	Conclusion	175

PART III: MEMORY-BASED NATURAL LANGUAGE UNDERSTANDING

Chapter 11	Using Memory in Text Understanding	
	<i>Michael Lebowitz</i>	181
	Abstract	181
	Introduction	181
	Basic <i>RESEARCHER</i> Understanding	
	Techniques	183
	Using Memory to Resolve Ambiguity	185
	Conclusion	192
 Chapter 12	 A More General Approach to Word Disambiguation	
	<i>Steven L. Lytinen</i>	195
	Abstract	195
	Introduction	195
	Previous Work in Semantics-Based	
	Word Disambiguation	196
	Problems With Lexically Based	
	Disambiguation Rules	199
	A Different Approach to Word	
	Disambiguation	203
	Conclusion	206
 Chapter 13	 Direct Memory Access Parsing	
	<i>Christopher K. Riesbeck and Charles E. Martin</i>	209
	Abstract	209
	Parsing Into Memory	209
	An Example System	210
	The Direct Memory Access Parser	212
	Recent DMAP Research	221
	Problems	223
	Conclusions	224

**Chapter 14 Parsing With Parallelism: A Spreading-Activation
Model of Inference Processing During Text
Understanding**

*Richard H. Granger, Kurt P. Eiselt, and
Jennifer K. Holbrook*

227

Abstract 227

Introduction 227

Background 228

The New Architecture 231

ATLAST: The Program 235

Summary and Conclusions 242

Appendix: Conditional Retention

Experiment 244

Author Index 247

Subject Index 251

1 Introduction

The chapters in this volume, though different in topic, style, and even scientific discipline, all share a common theme: that knowledge is first and foremost the product of experience. This obvious truth about human knowledge has significant ramifications for artificial intelligence (AI) and cognitive science. It means that to understand human knowledge and how it is organized and applied, we have to understand how it was acquired. It means that to claim that a representation of knowledge is adequate, we have to show how that knowledge could be learned and how it could be updated in the face of new experience.

KNOWLEDGE AND AI

The importance of knowledge for intelligence is one of AI's few consensual beliefs. Knowledge, in large quantities organized into usable chunks, is an essential ingredient of a great deal of intelligent behavior. How to chunk knowledge about the world to make it usable for reasoning is an unavoidable problem in AI research. The real world is full of objects and tableaux that have no identifying labels to help the visitor. To understand the world, we must carry around in our heads a field guide, a guide so large that the quality of its index is a major factor in the usefulness of the guide.

The metaphor is incomplete, however, until we make one further point: the guide is really a notebook that each of us writes and maintains. We are born with pencil and paper and have to take it from there. As a result, we each have a different field guide, in our own handwriting, with illegible scribbles, mistakes, corrections, and redundancies, and at no point can we stop and say that the book is complete.