



# PROBABILISTIC REASONING IN INTELLIGENT SYSTEMS:

Networks of Plausible Inference

*Judea Pearl*

*Department of Computer Science  
University of California  
Los Angeles*

**MORGAN KAUFMANN PUBLISHERS, INC.**  
**San Mateo, California**

Editor and President	Michael B. Morgan
Production Manager	Shirley Jowell
Text Design	Rick Van Genderen
Cover Design	Rick Van Genderen
Composition	Technically Speaking Publications
Copy Editor	Daniel Pearl
Proofreader	Linda Medoff

Grateful acknowledgement is made to the following for permission to reprint previously published material.

North-Holland, Amsterdam: Excerpts from Pearl, J., Fusion, Propagation, and Structuring in Belief Networks, *Artificial Intelligence* 29 (1986) 241-288; Pearl, J., Evidential Reasoning Using Stochastic Simulation of Causal Models, *Artificial Intelligence* 33 (1987) 173-215; Pearl, J., On Evidential Reasoning in a Hierarchy of Hypotheses, *Artificial Intelligence* 28 (1986) 9-15; and Pearl, J., Distributed Revision of Composite Beliefs, *Artificial Intelligence* 29 (1986) 241-288; Pearl, J., Embracing Causality in Default Reasoning, *Artificial Intelligence* 35 (1988) 259-271. Reprinted here with the permission of the publisher.

Pearl, J., Bayesian Decision Methods, *Encyclopedia of Artificial Intelligence*, Editor-in-Chief Stuart C. Shapiro, Vol. 1, pp 48-45 (1987). With permission of the publisher John Wiley & Sons.

### **Library of Congress Cataloging-in Publication Data is Available**

88-13069

ISBN 0-934613-73-7

Morgan Kaufmann Publishers, Inc. 2929 Campus Drive, Suite 260 San Mateo, CA 94403  
Order Fulfillment: P.O. Box 50490, Palo Alto, CA 94303 © 1988 by Morgan Kaufmann  
Publishers, Inc. All rights reserved. Printed in the United States of America

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 prior written permission of the publisher.

93 92 91 90 89    5 4 3 2 1

**In memory of my brother Beni**

# Preface

---

This book is the culmination of an investigation into the applicability of probabilistic methods to tasks requiring automated reasoning under uncertainty. The result is a computation-minded interpretation of probability theory, an interpretation that exposes the qualitative nature of this centuries-old formalism, its solid epistemological foundation, its compatibility with human intuition and, most importantly, its amenability to network representations and to parallel and distributed computation. From this vantage point I have attempted to provide a coherent account of probability as a language for reasoning with partial beliefs and bind it, in a unifying perspective, with other artificial intelligence (AI) approaches to uncertainty, such as the Dempster-Shafer formalism, truth maintenance systems, and nonmonotonic logic.

*Probabilistic Reasoning* has been written with a variety of readers in mind. It should be of interest to scholars and researchers in AI, decision theory, statistics, logic, philosophy, cognitive psychology, and the management sciences. Specifically, AI researchers can take an earnest look at probability theory, now that it is phrased in their language, and probabilists should be challenged by the new issues that emerge from the AI experiment. In addition, practitioners in the areas of knowledge-based systems, operations research, engineering, and statistics will find a variety of theoretical and computational tools that should be of immediate practical use. Application areas include diagnosis, forecasting, image understanding, multi-sensor fusion, decision support systems, plan recognition, planning and control, speech recognition – in short, almost any task requiring that conclusions be drawn from uncertain clues and incomplete information.

The book is also intended as a textbook for graduate-level courses in AI, operations research, and applied probability. In teaching this material at various levels of sophistication, I have found that the conceptual tools students acquire by treating the world probabilistically grow in value, even (perhaps especially) when the students go on to pursue other formalisms.

To my own surprise, most of these chapters turned out to be fairly self-contained, demanding only a basic understanding of the results established in previous chapters.

Chapter 1 identifies the basic AI paradigms of dealing with uncertainty and highlights the unique qualitative features that make probability theory a loyal guardian of plausible reasoning. Chapter 2 introduces the basic principles of

Bayesian inference and discusses some epistemological issues that emerge from this formalism. Those who have had no previous exposure to probability theory (some computer science students fall into this category) might wish to consult an introductory textbook and should follow closely the examples in Chapter 2 and work the exercises at the end of that chapter. In general, an elementary course in probability theory or decision analysis should be sufficient for mastering most of the book.

The casual reader seeking a painless glimpse at the basic issues of uncertainty should read the less technical sections in each chapter. These are indicated by a single asterisk (\*) in the Contents. Chapters 1, 9, and 10 will prove especially useful for those seeking a comprehensive look at how the various formalisms are related, and how they measure up to each other under the acid test of human intuition.

The more technically oriented reader will want to follow the sections marked with a double asterisk (\*\*), glancing occasionally at the definitions and results of other sections. This path leads the reader from traditional Bayesian inference and its graphical representations, into network propagation techniques (Chapters 4 and 5) and decision analysis (Chapter 6), and then into belief functions (Chapter 9) and default reasoning (Chapter 10). Knowledge engineers and developers of expert systems, for example, are advised to go straight to Section 3.3, then read Chapters 4, 5, 6, 7, and 9.

The most advanced sections, dealing with topics closer to current research frontiers, are marked with a triple asterisk (\*\*\*). These include the theory of graphoids (Chapter 3), learning methods (Chapter 8), and probabilistic semantics for default reasoning (Section 10.2).

The reader should not view these markings as strict delineators. Just as an advanced ski run has flat stretches and a beginner's run has a mogul or two, there will be occasional pointers to human-style reasoning in the midst of technical discussions, and references to computational issues in the midst of philosophical discussions. Some reviewers advised me to avoid this hybrid style of writing, but I felt that retreating toward a more traditional organization would deny the reader the sense of excitement that led me to these explorations. By confessing the speculative nature of my own curiosity I hope to encourage further research in these areas.

I owe a great debt to many people who assisted me with this work. First, I would like to thank the members of the Cognitive Systems Laboratory at UCLA, whose work and ideas formed the basis of many of these sections: Avi Dechter, Rina Dechter, Hector Geffner, Dan Geiger, Moises Goldszmidt, Jin Kim, Itay Meiri, Javier Pinto, Prasadram Ramachandra, George Rebane, Igor Roizen, Rony Ross, and Thomas Verma. Rina and Hector, in particular, are responsible for wresting me from the security blanket of probability theory into the cold darkness of constraint networks, belief functions, and nonmonotonic logic.

My academic and professional colleagues have been generous with their time and ideas. I have been most influenced by the ideas of Stephen Lauritzen, Glenn Shafer, and David Spiegelhalter, and my collaborations with Azaria Paz and Norman Dalkey have been very rewarding. Helpful comments on selected sections were offered by Ernst Adams, Moshe Ben-Bassat, Alan Bundy, Norman Dalkey, Johan de Kleer, Arthur Dempster, Richard Duda, David Heckerman, David Etherington, Max Henrion, Robert Goldman, Richard Jeffrey, Henry Kyburg, Vladimir Lifschitz, John Lowrance, David McAllester, John Pollock, Gregory Provan, Lenhart Schubert, Ross Shachter, Glenn Shafer and Michael Wellman.

The National Science Foundation deserves acknowledgment for sponsoring the research that led to these results, with special thanks to Y.T. Chien, Joe Dekens, and the late Ken Curtis, who encouraged my research when I was still a junior faculty member seeking an attentive ear. Other sponsors include Lashon Booker of the Navy Research Laboratory, Abraham Waksman of the Air-Force Office of Scientific Research, and Richard Weis of Aerojet Electro Systems.

The manuscript was most diligently typed, processed, illustrated and proofed by Gina George and Jackie Trang, assisted by Nina Roop and Lillian Casey. I thank the publisher for accommodating my idiosyncracies, and special thanks to a very special copy editor, Danny Pearl, whose uncompromising stance made these pages readable.

Finally, I owe a great debt to the rest of my family: to Tammy for reminding me why it all matters, to Michelle for making me feel useful and even drawing some of the figures, and especially to my wife Ruth for sheltering me from the travails of the real world and surrounding me with so much love, support, and hope.

**J.P.**

Los Angeles, California

June 1988

# Contents

---

## Chapter 1

<b>UNCERTAINTY IN AI SYSTEMS: AN OVERVIEW</b>	<b>1</b>
1.1 INTRODUCTION*	1
1.1.1 Why Bother with Uncertainty?	1
1.1.2 Why Is It a Problem?	2
1.1.3 Approaches to Uncertainty	2
1.1.4 Extensional vs. Intensional Approaches	3
1.2 EXTENSIONAL SYSTEMS: MERITS, DEFICIENCIES, AND REMEDIES*	4
1.2.1 Computational Merits	4
1.2.2 Semantic Deficiencies	6
1.2.3 Attempted Remedies and their Limitations	9
1.3 INTENSIONAL SYSTEMS AND NETWORK REPRESENTATIONS*	12
1.3.1 Why Networks?	12
1.3.2 Graphoids and the Formalization of Relevance and Causality	13
1.4 THE CASE FOR PROBABILITIES*	14
1.4.1 Why Should Beliefs Combine Like Frequencies?	15
1.4.2 The Primitive Relationships of Probability Language	16
1.4.3 Probability as a Faithful Guardian of Common Sense	19

---

\* Basic Issues

\*\* Technical and Substantive Discussions

\*\*\* Advanced Research

For a detailed explanation of these levels, see Preface.



1.5	QUALITATIVE REASONING WITH PROBABILITIES*	23
1.5.1	Softened Logic vs. Hardened Probabilities	24
1.5.2	Probabilities and the Logic of “Almost True”	25
1.6	BIBLIOGRAPHICAL AND HISTORICAL REMARKS	26

*Chapter 2*

**BAYESIAN INFERENCE** 29

2.1	BASIC CONCEPTS	29
2.1.1	Probabilistic Formulation and Bayesian Inversion*	29
2.1.2	Combining Predictive and Diagnostic Supports*	34
2.1.3	Pooling of Evidence*	36
2.1.4	Recursive Bayesian Updating*	37
2.1.5	Multi-Valued Hypotheses**	39
2.2	HIERARCHICAL MODELING	42
2.2.1	Uncertain Evidence (Cascaded Inference)*	42
2.2.2	Virtual (Intangible) Evidence*	44
2.2.3	Predicting Future Events**	47
2.2.4	Multiple Causes and “Explaining Away”*	49
2.2.5	Belief Networks and the Role of Causality*	50
2.3	EPISTEMOLOGICAL ISSUES OF BELIEF UPDATING	52
2.3.1	Patterns of Plausible Inference: Polya vs. Bayes?*	52
2.3.2	The Three Prisoners Paradox: When the Bare Facts Won’t Do*	58
2.3.3	Jeffrey’s Rule and the Problem of Autonomous Inference Agents*	62
2.4	BIBLIOGRAPHICAL AND HISTORICAL REMARKS	70

EXERCISES 73

*Chapter 3*

**MARKOV AND BAYESIAN NETWORKS: Two Graphical Representations of Probabilistic Knowledge** 77

3.1	FROM NUMERICAL TO GRAPHICAL REPRESENTATIONS	78
3.1.1	Introduction*	78

3.1.2	An Axiomatic Basis for Probabilistic Dependencies**	82
3.1.3	On Representing Dependencies by Undirected Graphs*	90
3.1.4	Axiomatic Characterization of Graph-Isomorph Dependencies***	93
3.2	MARKOV NETWORKS	96
3.2.1	Definitions and Formal Properties***	96
3.2.2	Illustrations**	100
3.2.3	Markov Network as a Knowledge Base***	104
3.2.4	Decomposable Models**	108
3.3	BAYESIAN NETWORKS	116
3.3.1	Dependence Semantics for Bayesian Networks**	117
3.3.2	Bayesian Network as a Knowledge Base**	122
3.3.3	How Expressive Are Bayesian Networks?***	126
3.4	BIBLIOGRAPHICAL AND HISTORICAL REMARKS	131
	EXERCISES	134
	APPENDIX 3-A Proof of Theorem 3	139
	APPENDIX 3-B Proof of Theorem 4	141
 <i>Chapter 4</i>		
	<b>BELIEF UPDATING BY NETWORK PROPAGATION</b>	143
4.1	AUTONOMOUS PROPAGATION AS A COMPUTATIONAL PARADIGM	144
4.1.1	Constraint Propagation and Rule-based Computation*	145
4.1.2	Why Probabilistic Reasoning Seems to Resist Propagation*	148
4.2	BELIEF PROPAGATION IN CAUSAL TREES	150
4.2.1	Notation**	150
4.2.2	Propagation in Chains**	153
4.2.3	Propagation in Trees**	162

4.3	BELIEF PROPAGATION IN CAUSAL POLYTREES (SINGLY CONNECTED NETWORKS)	175
4.3.1	Propagation Rules**	177
4.3.2	Canonical Models of Multicausal Interactions**	184
4.4	COPING WITH LOOPS	195
4.4.1	Clustering Methods*	197
4.4.2	The Method of Conditioning (Reasoning by Assumptions)*	204
4.4.3	Stochastic Simulation**	210
4.5	WHAT ELSE CAN BAYESIAN NETWORKS COMPUTE?	223
4.5.1	Answering Queries*	223
4.5.2	Introducing Constraints*	225
4.5.3	Answering Conjunctive Queries**	226
4.6	BIBLIOGRAPHICAL AND HISTORICAL REMARKS	232
	EXERCISES	234
	APPENDIX 4-A Auxilliary Derivations for Section 4.5.3	236
<i>Chapter 5</i>		
	<b>DISTRIBUTED REVISION OF COMPOSITE BELIEFS</b>	239
5.1	INTRODUCTION	239
5.2	ILLUSTRATING THE PROPAGATION SCHEME	241
5.2.1	Belief Updating (A Review)*	242
5.2.2	Belief Revision*	245
5.3	BELIEF REVISION IN SINGLY CONNECTED NETWORKS	250
5.3.1	Deriving the Propagation Rules***	251
5.3.2	Summary of Propagation Rules*	255
5.3.3	Reaching Equilibrium and Assembling a Composite Solution**	257
5.3.4	Comparison to Belief Updating*	259
5.3.5	Coping with Loops*	260

5.4	DIAGNOSIS OF SYSTEMS WITH MULTIPLE FAULTS	263
5.4.1	An Example: Electronic Circuit**	263
5.4.2	Initialization**	265
5.4.3	Fault Interpretation**	267
5.4.4	Finding the Second-Best Interpretation**	271
5.5	APPLICATION TO MEDICAL DIAGNOSIS	272
5.5.1	The Causal Model*	273
5.5.2	Message Propagation**	275
5.5.3	Choosing the Best Interpretation*	278
5.5.4	Generating Explanations*	279
5.5.5	Reversibility vs. Perseverance*	280
5.6	THE NATURE OF EXPLANATIONS	281
5.6.1	Accepting vs. Assessing Beliefs*	281
5.6.2	Is a Most-Probable Explanation Adequate?*	283
5.6.3	Circumscribing Explanations*	285
5.7	CONCLUSIONS	286
5.8	BIBLIOGRAPHICAL AND HISTORICAL REMARKS	287
	EXERCISES	288
<i>Chapter 6</i>		
	<b>DECISION AND CONTROL</b>	289
6.1	FROM BELIEFS TO ACTIONS: INTRODUCTION TO DECISION ANALYSIS	289
6.1.1	Rational Decisions and Quality Guarantees*	289
6.1.2	Consequences, Payoffs, and Lotteries*	290
6.1.3	Calibrating the Value of a Lottery*	292
6.1.4	The Axioms of Utility Theory**	294
6.1.5	Utility Functions*	297
6.2	DECISION TREES AND INFLUENCE DIAGRAMS	299
6.2.1	Decision Trees*	299
6.2.2	Planning with Decision Trees*	300
6.2.3	Influence Diagrams*	306

6.3	THE VALUE OF INFORMATION	313
6.3.1	Information Sources and Their Values**	313
6.3.2	Myopic Assessments of Information Sources**	314
6.4	RELEVANCE-BASED CONTROL	318
6.4.1	Focusing Attention**	318
6.4.2	Utility-Free Assessment of Information Sources***	319
6.4.3	Controlling Attention*	324
6.4.4	Summary*	326
6.5	BIBLIOGRAPHICAL AND HISTORICAL REMARKS	327
	EXERCISES	328
 <i>Chapter 7</i>		
	<b>TAXONOMIC HIERARCHIES, CONTINUOUS VARIABLES, AND UNCERTAIN PROBABILITIES</b>	333
7.1	EVIDENTIAL REASONING IN TAXONOMIC HIERARCHIES	333
7.1.1	Taxonomic vs. Causal Hierarchies*	333
7.1.2	Evidence Propagation in Taxonomic Hierarchies*	337
7.1.3	Probabilistic Justification**	340
7.1.4	Psychological and Computational Characteristics*	342
7.2	MANAGING CONTINUOUS VARIABLES	344
7.2.1	Plausible Reasoning about Uncertain Quantities*	344
7.2.2	Propagating Estimates and Ranges**	346
7.2.3	Qualitative Patterns of Reasoning**	351
7.2.4	Discussion*	355
7.3	REPRESENTING UNCERTAINTY ABOUT PROBABILITIES	357
7.3.1	The Semantics of Probabilities of Probabilities*	358
7.3.2	De Finetti's Paradigm of Uncertain Contingencies*	359
7.3.3	A Formal Definition of Network-Induced Confidence Measures**	363

7.3.4	The Effect of Evidence on Confidence: An Example**	366
7.3.5	Conclusions*	372
7.4	BIBLIOGRAPHICAL AND HISTORICAL REMARKS	372
	EXERCISES	374
	APPENDIX 7-A Derivation of Propagation Rules For Continuous Variables	375
 <i>Chapter 8</i>		
	<b>LEARNING STRUCTURE FROM DATA</b>	381
8.1	CAUSALITY, MODULARITY, AND TREE STRUCTURES*	383
8.2	STRUCTURING THE OBSERVABLES	387
8.2.1	Chow's Method of Constructing Trees**	387
8.2.2	Structuring Causal Polytrees**	390
8.2.3	Conclusions: Causation or Covariation?*	396
8.3	LEARNING HIDDEN CAUSES	398
8.3.1	Problem Definition and Nomenclature***	398
8.3.2	Star-Decomposable Triplets***	401
8.3.3	A Tree-Reconstruction Procedure***	402
8.3.4	Extensions to Normal Variables***	405
8.3.5	Conclusions and Open Questions*	407
8.4	BIBLIOGRAPHICAL AND HISTORICAL REMARKS	408
	EXERCISES	409
	APPENDIX 8-A Proof of Theorems 1 and 2	411
	APPENDIX 8-B Conditions for Star-Decomposability ( <i>After Lazarfeld [1966]</i> )	412
 <i>Chapter 9</i>		
	<b>NON-BAYESIAN FORMALISMS FOR MANAGING UNCERTAINTY</b>	415
9.1	THE DEMPSTER-SHAFER THEORY	416
9.1.1	Basic Concepts*	416

9.1.2	Comparing Bayesian and Dempster-Shafer Formalisms*	421
9.1.3	Dempster's Rule of Combination**	428
9.1.4	More on the Nature of Probability Intervals*	438
9.1.5	Applications to Rule-based Systems**	440
9.1.6	Bayes vs. Dempster-Shafer: A Semantic Clash*	447
9.2	TRUTH MAINTENANCE SYSTEMS	450
9.2.1	Naming the Assumptions*	451
9.2.2	Uncertainty Management in an ATMS**	454
9.2.3	Incidence Calculus*	456
9.3	PROBABILISTIC LOGIC**	457
9.4	BIBLIOGRAPHICAL AND HISTORICAL REMARKS	462
	EXERCISES	465
 <i>Chapter 10</i>		
	<b>LOGIC AND PROBABILITY: THE STRANGE CONNECTION</b>	467
10.1	INTRODUCTION TO NONMONOTONIC REASONING*	467
10.1.1	Reiter's Default Logic*	468
10.1.2	Problems with Default Logics*	470
10.1.3	Empirical vs. Procedural Semantics in Default Reasoning*	473
10.2	PROBABILISTIC SEMANTICS FOR DEFAULT REASONING	481
10.2.1	$\epsilon$ -semantics***	482
10.2.2	Axiomatic Characterization and a System of Defeasible Inference***	485
10.2.3	Relevance-based Conventions***	491
10.2.4	Do People Use the Logic of Extreme Probabilities?*	493

10.3	EMBRACING CAUSALITY IN DEFAULT REASONING	497
10.3.1	How Old Beliefs Were Established Determines Which New Beliefs Are Evoked*	497
10.3.2	More on the Distinction Between Causal and Evidential Support*	499
10.3.3	The C-E System: A Coarse Logical Abstraction of Causal Directionality*	501
10.3.4	Implicit Suppressors and the Need for Finer Abstractions*	507
10.4	A PROBABILISTIC TREATMENT OF THE YALE SHOOTING PROBLEM	509
10.4.1	The Problem and its Solution*	509
10.4.2	Concluding Dialogue*	513
10.5	BIBLIOGRAPHICAL AND HISTORICAL REMARKS	516
	EXERCISES	518
	<b>Bibliography</b>	521
	<b>Author Index</b>	539
	<b>Subject Index</b>	545



## Chapter 1

# UNCERTAINTY IN AI SYSTEMS: AN OVERVIEW

---

*I consider the word probability as meaning the state of mind with respect to an assertion, a coming event, or any other matter on which absolute knowledge does not exist.*

— August De Morgan, 1838

## 1.1 INTRODUCTION

### 1.1.1 Why Bother with Uncertainty?

Reasoning about any realistic domain always requires that some simplifications be made. The very act of preparing knowledge to support reasoning requires that we leave many facts unknown, unsaid, or crudely summarized. For example, if we choose to encode knowledge and behavior in rules such as "Birds fly" or "Smoke suggests fire," the rules will have many exceptions which we cannot afford to enumerate, and the conditions under which the rules apply (e.g., seeing a bird or smelling smoke) are usually ambiguously defined or difficult to satisfy precisely in real life. Reasoning with exceptions is like navigating a minefield: Most steps are safe, but some can be devastating. If we know their location, we can avoid or defuse each mine, but suppose we start our journey with a map the size of a postcard, with no room to mark down the exact location of every mine or the way they are wired together. An alternative to the extremes of ignoring or enumerating exceptions is to *summarize* them, i.e., provide some warning signs to indicate which areas of the minefield are more dangerous than others. Summarization is