Uncertainty in Artificial Intelligence

Edited by

Laveen N. KANAL

and

John F. LEMMER

Uncertainty in Artificial Intelligence

Edited by

Laveen N. KANAL
University of Maryland
College Park
Maryland
U.S.A.

and

John F. LEMMER
Knowledge Systems Concepts
Rome
New York
U.S.A.

1986

NORTH-HOLLAND
AMSTERDAM · NEW YORK · OXFORD · TOKYO

ISBN: 0 444 70058 7

Publishers:
ELSEVIER SCIENCE PUBLISHERS B.V.
P.O. Box 1991
1000 BZ Amsterdam
The Netherlands

Sole distributors for the U.S.A. and Canada:
ELSEVIER SCIENCE PUBLISHING COMPANY, INC.
52 Vanderbilt Avenue
New York, N.Y. 10017
U.S.A.

Library of Congress Cataloging-in-Publication Data

Uncertainty in artificial intelligence.

(Machine intelligence and pattern recognition ; v. 4)
Bibliography: p.
1. Artificial intelligence. 2. Uncertainty
(Information theory) I. Kanal, Laveen N. II. Lemmer,
John F. III. Series.
Q335.U53 1986 006.3 86-13408
ISBN 0-444-70058-7

PRINTED IN THE NETHERLANDS

UNCERTAINTY IN ARTIFICIAL INTELLIGENCE

Machine Intelligence and Pattern Recognition

Volume 4

Series Editors

L. N. KANAL

and

A. ROSENFELD

University of Maryland
College Park
Maryland
U.S.A.

NORTH-HOLLAND
AMSTERDAM · NEW YORK · OXFORD · TOKYO

PREFACE

This book had its origins in the panel discussion on uncertainty held during the 1984 AAAI Conference in Austin, Texas. At that time, it became clear that many people had many different ideas on how to handle uncertainty in Artificial Intelligence. One of us (J.F.L.) became interested in promoting further discussion among the proponents of the various different views. This interest, along with Peter Cheeseman's collaboration, resulted in the Workshop on Probability and Uncertainty in AI held in conjunction with the 1985 International Joint Conference on AI in Los Angeles. Several excellent papers were presented at this workshop and many ideas were exchanged. The current volume contains some papers as originally presented at the workshop, revised versions of a number of the papers presented at the workshop, and additional papers written especially for this volume. In most cases, revisions were made in the light of workshop discussions. The result is a volume which presents many diverse opinions about handling uncertainty in AI and which makes it evident that there is still much to argue about on this topic.

This volume brings together a wide range of perspectives on uncertainty, often by their principal proponents. The first section consists of two introductory papers. The first paper reviews numeric and non-numeric approaches currently being proposed for handling uncertainty in AI systems and comments on a number of the perspectives presented in this volume. The second reviews consensus rules for combining experts' opinions. The next set of papers is devoted to explications or critiques of current approaches to uncertainty. This is followed by a set of papers presenting a synthesis of current approaches; a set of papers describing architectures and algorithms for building systems which incorporate uncertainty; and a set of papers addressing techniques for inducing uncertain information. The book concludes with some papers offering alternative perspectives on uncertainty in AI and on using models of uncertainty to address problems of minimax game trees.

Some of the notable issues which emerge from reading the papers in this volume are the following: It is surprising to note that despite the widespread enthusiasm for an interval based calculus of uncertainty, as yet there appears to be no decision theory based on intervals. Nor does Dempster—Shafer theory appear adequate for empirical data. Also surprising is that many authors interested in AI are now advocating probability as the best numeric model for uncertainty when only a few years ago probability was thought by many in AI to be inadequate. Arguments are even being advanced that people are Bayesian after all. Nevertheless there remain strong dissenting opinions not only about probability but even about the utility of any numeric method in this context.

We would like to thank all those who helped in organizing the workshop and all the authors for their enthusiastic collaboration; without them this volume would not exist. It is our hope that this work will stimulate further communication among workers seriously concerned with handling uncertainty in AI systems which deal with real world problems.

Laveen N. Kanal
College Park, MD

John F. Lemmer
Rome, NY

CONTRIBUTORS

Bruce Abramson *Columbia University, NY 10027 and University of California, Los Angeles, CA 90024*

Malcolm Bauer *Bell Communications Research, Morristown, NJ 07974*

Carlos Berenstein *University of Maryland and L.N.K. Corporation, College Park, MD 20742*

Raj K. Bhatnagar *University of Maryland, College Park, MD 20742*

Piero P. Bonissone *General Electric Corporate Research and Development, Schenectady, NY 12301*

Jack Breese *Stanford University, Stanford, CA 94305*

B. Chandrasekaran *Ohio State University, Columbus, OH 43210*

P. Cheeseman *NASA Ames Research Center, Moffett Field, CA 94035*

Chee-Yee Chong *Advanced Decision Systems, Mountain View, CA 94040*

Marvin S. Cohen *Decision Science Consortium, Inc., Falls Church, VA 22043*

N.C. Dalkey *University of California, Los Angeles, CA 90024*

Keith S. Decker *General Electric Corporate Research and Development, Schenectady, NY 12301*

John Fox *Imperial Cancer Research Fund Laboratories, London WC2A 3PX*

Robert M. Fung *Advanced Decision Systems, Mountain View, CA 94040*

Matthew L. Ginsberg *Stanford University, Stanford, CA 94305*

Benjamin N. Grosof *Stanford University, Stanford, CA 94305*

Henry Hamburger *George Mason University, Fairfax, VA 22030 and Navy Center for Applied Research in Artificial Intelligence*

Stephen Jose Hanson *Bell Communications Research, Morristown, NJ 07974*

David Heckerman *Stanford University, Stanford, CA 94305*

Max Henrion *Carnegie-Mellon University, Pittsburgh, PA 15213*

Peter D. Holden *McDonnell Douglas Corporation, St. Louis, MO 63166*

Samuel Holtzman *Strategic Decisions Group, Menlo Park, CA 94025*

Eric Horvitz *Stanford University, Stanford, CA 94305*

Daniel Hunter *TRW/DSG, Redondo Beach, CA 90278*
Rodney W. Johnson *Naval Research Laboratory, Washington, DC 20375*
Laveen N. Kanal *University of Maryland and L.N.K. Corporation, College Park, MD 20742*
Karl G. Kempf *McDonnell Douglas Corporation, St. Louis, MO 63166*
David Lavine *L.N.K. Corporation, College Park, MD 20742*
John F. Lemmer *Knowledge Systems Concepts, Inc., Rome, NY 13440*
Todd S. Levitt *Advanced Decisions Systems, Mountain View, CA 94040*
Gerald Shao-Hung Liu *Schlumberger/Sentry, San Jose, CA 95115*
Ronald P. Loui *University of Rochester, Rochester, NY 14627*
Dana Nau *University of Maryland, College Park, MD 20742*
Judea Pearl *University of California, Los Angeles, CA 90024*
Bruce M. Perrin *McDonnell Douglas Corporation, St. Louis, MO 63166*
Paul Purdom *Indiana University, Bloomington, IN 47405*
Larry Rendell *University of Illinois at Urbana-Champaign, IL 61801*
Ross D. Shachter *Stanford University, Stanford, CA 94305*
Glenn Shafer *University of Kansas, Lawrence, KS 66045*
John E. Shore *Naval Research Laboratory, Washington, DC 20375*
Ray Solomonoff *Oxbridge Research, Cambridge, MA 02238*
David J. Spiegelhalter *MRC Biostatistics Unit, Cambridge, England*
Michael C. Tanner *Ohio State University, Columbus, OH 43210*
Chun-Hung Tzeng *Ball State University, Muncie, IN 47306*
David S. Vaughn *McDonnell Douglas Corporation, St. Louis, MO 63166*
Ben P. Wise *Carnegie-Mellon University, Pittsburgh, PA 15213*
Robert M. Yadrick *McDonnell Douglas Corporation, St. Louis, MO 63166*
Ronald R. Yager *Iona College, New Rochelle, NY 10801*
Lotfi A. Zadeh *University of California, Berkely, CA 94720*
Alf C. Zimmer *University of Regensburg, F.R.G.*

CONTENTS

III. SYNTHESIS OF CURRENT APPROACHES TO UNCERTAINTY

IV. INCORPORATING UNCERTAINTY IN SYSTEMS

I

OVERVIEWS AND REVIEWS

Uncertainty in Artificial Intelligence
L.N. Kanal and J.F. Lemmer (Editors)

Handling Uncertain Information : A Review of Numeric and Non-numeric Methods

Raj K. Bhatnagar and Laveen N. Kanal

Machine Intelligence and Pattern Analysis Laboratory,
Department of Computer Science,
University of Maryland,
College Park, MD 20742, USA.

Problem solving and decision making by humans is often done in environments where information concerning the problem is partial or approximate. AI researchers have been attempting to emulate this capability in computer expert systems. Most of the methods used to-date lack a theoretical foundation. Some theories for handling uncertainty of information have been proposed in the recent past. In this paper, we critically review these theories. The main theories that we examine are: Probability Theory, Shafer's Evidence Theory, Zadeh's Possibility Theory, Cohen's Theory of Endorsements and the non-monotonic logics. We describe these in terms of the representation of uncertain information, and combination of bodies of information and inferencing with such information, and consider the strong and weak aspects of each theory.

1. *Introduction*

In the recent past, one focus of research in Artificial Intelligence has been the problem of " Approximate Reasoning ". The problem deals with decision making and reasoning processes in situations where the information is deficient in one or more of the following ways: information is partial, the information is not fully reliable, the representation language is inherently imprecise and information from multiple sources is conflicting. All these uncertainties may exist for the knowledge relating to the occurence of events in a world model and also for causal or any other relationships among various events. Information is partial when answers to some relevant questions are not known. It is said to be approximate when the answers are known but are not accurate and exact. It may be caused either because of a partially reliable source or imprecise representation language. In the numeric context, approximation can be viewed as a value with a known error margin. When the possible values for a variable are symbolic rather than numeric, approximations can be represented in terms of a fuzzy set with a corresponding membership function.

Causal relationships between any two events in a real world situation are not always easy to specify. Any event seems to be affected or related to numerous other factors in one's world model. A rule in an expert system then, would require too many antecedents for a consequent. One solution is to have a very restricted model of the world but that may not be a very useful option. An easier

This work has been supported in part by grants from the National Science Foundation to the Machine Intelligence and Pattern Analysis Laboratory.

way out is to specify heuristic rules along with some measure of the confidence one has in the relationship specified by the rule. Such rules add to the uncertainty of the inferences obtained from premises which are uncertain. If the information concerning a problem is partial or approximate, then the problem can only be solved approximately, *i.e.* , with uncertainty. A solution without uncertainty may be possible when complete and exact information is available. But in this case also, we may settle for an approximate solution in order to reduce the computational cost. The distributed environments for problem solving may also contribute to the uncertainty because the information flow among nodes may have to be limited. A detailed discussion on theoretical aspects of uncertainty in problem solving environments can be found in Traub [22].

We examine various approaches that have been suggested for solving problems in environments in which the information is uncertain. Three different aspects of this problem are:

-Representation of uncertain information

-Combination of bodies of uncertain information

-Drawing of inferences using uncertain information

We consider these problems one by one as they are handled by various numeric and non-numeric methods proposed in the recent literature. In particular, we examine probability Theory, Shafer's Evidence Theory, Zadeh's Possibility Theory, A Theory of Endorsements and the non-monotonic logics.

2. *Representation Of Uncertain Information*

On first thought the problem of representing uncertain information itself appears to be very vague. It is because an inexact concept can be made clear only by providing informal explanations, examples and exceptions etc. Any formal representation consisting of only a finite number of symbols, each having a predefined and exact interpretation can represent only an exact idea. To be useful, any such representation must fulfil the following requirements: (1) Similarity to the actual concept to be explained and (2) simplicity. Intuitively, one might argue that doing away with the exactness, that is, allowing inexact interpretations for syntactic entities, might help. But we can't do meaningful symbolic processing with such representations. Any formalism, to be amenable to mathematical treatment must be exact; yet the problem is to represent an incompletely known world model.

When a decision making system wants to select an alternative from amongst many, it should know the state of the existing world. An alternative may have different payoffs in different states of the world. Therefore, state information is necessary in order to be able to select the best alternative. It is this information about the state of the world, which in most cases is uncertain. If extensive expert human interaction is part of a decision making system, then one may agree with Chandrasekaran[4] that resolution of uncertainty should be left to the human who is an expert in that domain. Chandrasekaran also states that humans do not use a single method for resolving uncertainties of various types and that in attempts to understand intelligent human problem solving strategies using programs, a " search for normative uncertainty calculi is pointless ". It is easy to