UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE

AD-A232 140

DTIC FILE COPY

REPORT DOCUMENTATION PAGE

DTIC ELECTE FEB 25 1991 S D D

1a. REPORT SECURITY CLASSIFICATION Unclassified	1b. RESTRICTIVE MARKINGS
2a. SECURITY CLASSIFICATION AUTHORITY	3. DISTRIBUTION/AVAILABILITY OF REPORT Approved for public release; distribution unlimited.
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE	
4. PERFORMING ORGANIZATION REPORT NUMBER(S) ECU-8139	5. MONITORING ORGANIZATION REPORT NUMBER(S) ARO 26314.2-EL
6a. NAME OF PERFORMING ORGANIZATION SRI International / 6b. OFFICE SYMBOL (If applicable)	7a. NAME OF MONITORING ORGANIZATION U. S. Army Research Office
6c. ADDRESS (City, State, and ZIP Code) 333 Ravenswood Avenue, Menlo Park, CA 94025	7b. ADDRESS (City, State, and ZIP Code) P. O. Box 12211, Research Triangle Park, NC 27709-2211
8a. NAME OF FUNDING/SPONSORING ORGANIZATION U. S. Army Research Office / 8b. OFFICE SYMBOL (If applicable)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER DAAL03-89-K-0156
8c. ADDRESS (City, State, and ZIP Code) P. O. Box 12211, Research Triangle Park, NC 27709-2211	10. SOURCE OF FUNDING NUMBERS: PROGRAM ELEMENT NO / PROJECT NO. / TASK NO. / WORK UNIT ACCESSION NO

11 TITLE (Include Security Classification)

ADVANCED METHODS OF APPROXIMATE REASONING (Unclassified)

12 PERSONAL AUTHOR(S)

Dr. Enrique H. Ruspini

13a. TYPE OF REPORT	13b. TIME COVERED	14. DATE OF REPORT (Year, Month, Day)	15. PAGE COUNT
Final	FROM 9/89 TO 9/90	November 30, 1990	120 pages

16. SUPPLEMENTARY NOTATION The view, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy, or decision, unless so designated by other documentation.

17 COSATI CODES			18 SUBJECT TERMS (Continue on reverse if necessary and identify by block number)	
FIELD	GROUP	SUB-GROUP	Approximate Reasoning	Probabilistic Reasoning
			Uncertainty	Fuzzy Logic
			Imprecision	Possibilistic Reasoning

19 ABSTRACT (Continue on reverse if necessary and identify by block number)

This research was directed toward establishing basic conceptual foundations for approximate reasoning concepts and toward the development of frameworks that facilitate the development and comparison of applicable techniques. Approximate reasoning is the common name utilized to to describe automated techniques for the representation and manipulation of imprecise, uncertain, unreliable, and vague information.

Our attention was focused primarily on the development of conceptual bases for possibilistic or "fuzzy" logic. Using a conceptual framework, previously employed to explain the meaning of the Dempster-Shafer calculus of evidence (i.e., possible-world semantics), we developed a semantic model for possibilistic logics that clearly shows them to be substantially different from their probabilistic counterparts both in meaning and in formal structure.

20. DISTRIBUTION/AVAILABILITY OF ABSTRACT ☐ UNCLASSIFIED/UNLIMITED ☐ SAME AS RPT. ☐ DTIC USERS	21. ABSTRACT SECURITY CLASSIFICATION Unclassified	
22a NAME OF RESPONSIBLE INDIVIDUAL Enrique H. Ruspini	22b. TELEPHONE (Include Area Code) (415)859-2314	22c. OFFICE SYMBOL AI Center

DD FORM 1473, 84 MAR 83 APR edition may be used until exhausted. All other editions are obsolete.

SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED

91 2 15 046

18. (continuation)

Evidential Reasoning
Dempster-Shafer Theory
Artificial Intelligence

19. ABSTRACT (continuation)

At the semantic level, possibilistic techniques were shown to be based on the notion of similarity, unlike probabilistic techniques, which seek to determine the likelihood of certain statements as true descriptors of the state of affairs in the real world. At a purely formal level, possibilistic techniques were shown to be the result of imposing metric structures upon a set of possible worlds rather than the consequence of defining set measures (i.e., probabilities) on that set.

The model developed during this research has led to better characterizations of the basic structures of possibilistic logic (notably, the notion of conditional possibility) and its basic rules of derivation (i.e., the generalized modus ponens). Furthermore, we have also clarified the role of the notions of possibility and utility as potential sources of similarity measures.

We report also on related work that extends the scope of the Dempster-Shafer calculus of evidence while clarifying certain conceptual misunderstandings about its generality and applicability.

The research results reported herein clarify fundamental aspects of information processing under conditions of imprecision and uncertainty, being particularly relevant to the development of analogical reasoning systems, i.e., automated devices that exploit similarities between scenarios to "extrapolate" from known examples into unknown situations. Because of their fundamental nature, these results are applicable to a wide variety of problems of Army interest including intelligence analysis, intelligent device and system control, vulnerability analysis, human-factors engineering, material analysis, fault diagnosis, reliability analysis, system design, and mission planning and counterplanning.

ADVANCED CONCEPTS OF APPROXIMATE REASONING

FINAL TECHNICAL REPORT

Enrique H. Ruspini

November 30, 1990

Prepared for

U.S. ARMY RESEARCH OFFICE

Contract No. DAAL03-89-K-0156

Artificial Intelligence Center
SRI International
Menlo Park, California

ADVANCED CONCEPTS OF APPROXIMATE REASONING

FINAL TECHNICAL REPORT

Executive Summary

Enrique H. Ruspini
Artificial Intelligence Center
SRI International

1 Introduction

This final report consists primarily of a collection of papers that have been published, presented, or await publication in various forums presenting the results of research, sponsored by the U.S. Army Research Office, on an artificial intelligence discipline known as "Approximate Reasoning."

This collection includes both detailed technical presentations of approximate reasoning issues [1,6], various summaries of those presentations [4,5,8], and an encompassing overview of their significance [2] in the context of a unified formal framework, developed as part of the reported research.

For this reason, we have chosen a format based on inclusion of all papers relevant to our research, preceded by this executive summary, which is also intended to guide the interested reader to the diverse works that make the bulk of the report.

The research program on advanced concepts of approximate reasoning had the goal of establishing firm formal foundations that explain the different technologies proposed to solve the problems associated with the processing of imprecise and uncertain information, permit a comparison of their advantages and disadvantages, and, specially, allow the determination of their applicability to specific problems.

The research results reported herein clarify fundamental aspects of information processing under conditions of imprecision and uncertainty. These results represent particularly important steps toward the development of systems for *analogical reasoning*, i.e., automated devices that exploit similarities between scenarios to "extrapolate" from known examples into unknown situations.

Because of their fundamental nature, these results are applicable to a wide variety of problems of Army interest including intelligence analysis, autonomous device planning and control, vulnerability analysis, human factors engineering, material analysis, fault diagnosis, reliability analysis, system design, and mission planning and counterplanning.

On the basis of the nature of the results obtained during this research and current practical experience with the applicability of various approximate reasoning techniques, it is possible to identify the following applications as being particularly amenable to treatment in the near future:

1. *Control of unstable systems*, such as helicopters, land vehicles, or weapon platforms, by means of possibilistic control techniques
2. *Control of navigation*, *target tracking*, and *obstacle avoidance* by autonomous mobile agents
3. *Elimination of involuntary platform/hand movement* in object-tracking tasks.
4. *Development of vulnerability measures* and related assessments of structural viability.
5. Development of *approximate models* of complex systems.
6. *Coordination of real time intelligent agents* on the basis of considerations about their usefulness, associated risks, and probability of success.

2 Approximate Reasoning

Approximate Reasoning is the collective name given to a variety of automated methods and techniques for the analysis of imprecise and uncertain information.

The first task in our investigation was to clarify the nature of the approximate reasoning problem: a poorly understood question that was felt to be the basic cause of the controversy that characterized the state of the art. Prior characterizations of approximate reasoning technology broadly interpreted the epithet "approximate" as an indication of either the poor quality of the underlying knowledge or that of the proposed techniques, considered to be heuristical imitations of the sounder methods of classical logic.

Our approach to the characterization of the approximate-reasoning problem was based on continuation of previous work of the principal investigator ("The Logical Foundations of Evidential Reasoning," SRI AIC Technical Note No. 408, 1987), which relied on the logical notion of "possible world." The result of these investigations was the development of a unified framework for the approximate reasoning problem that is briefly summarized in a paper presented at the Fourth International Symposium on Knowledge and its Engineering [7] and that is considerably expanded in a related assessment of the state of the art and its progress [2].

Informally speaking, possible worlds are the conceivable situations, scenarios, states, or behaviors of a real-world system, i.e., the conceivable solutions of a typical situation- or state-assessment problem. In those problems, we are typically required to state whether the system in question (e.g., "the weather at Menlo Park") is (or was, or will be) in such a state that certain statements (called hypotheses) about it are true (e.g., "... will be rainy on November 15").

To answer such questions in the context of a typical reasoning problem, we usually make various observations of our system (e.g., temperatures, pressures) that, when combined with existing background knowledge (e.g., meteorology), eliminate certain conceivable possibilities from consideration. The remaining states, called in our model the *evidential set* because of its obvious relationship with observed evidence, are then examined to determine whether all its possible states are such that a hypothesis of interest is true in all of them or it is false in all of them.

If that is indeed the case, as illustrated in Figure 1, then the problem is a conventional reasoning problem capable of being, at least conceptually, solved by classical logical techniques (i.e., the *evidence* implies the *hypothesis*.).

In an approximate reasoning problem, however, the situation resembles that illustrated in Figure 2, where, in some of the possibilities that are consistent with the evidence are such that the hypothesis is true, while on others it is false. Being faced with such an inability to solve the problem of finding whether a hypothesis is true or false, all approximate reasoning methods, in one way or another, modify the problem to be solved concentrating instead in describing the evidential set in terms of its relationship with the hypothesis of interest.

Probabilistic reasoning methods, illustrated in Figure 3, for example, seek to determine the proportion of evidential possibilities where a hypothesis is true (i.e., the conditional probability of truth). This proportion is usually estimated with the aid of statistical tables that summarize experience under similar circumstances.

Possibilistic reasoning methods, on the other hand, rely on measures of resemblance and similarity to determine, as illustrated in Figure 4, to what extent evidential possibilities resemble, or are close to, the set of possibilities where the hypothesis is true. The similarity measure that makes such a characterization possible is intended to be a measure of the extent by which facts that are true in one situation or scenario are true in another. For example, assessments of the stability of a weapon platform under some assumptions will remain approximately valid for similar platforms.

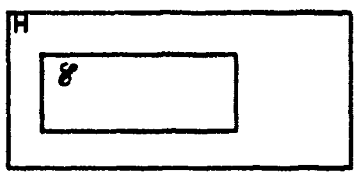

Figure 1: The conventional reasoning problem.

3 Possibilistic Reasoning

Having in the past successfully utilized possible-world models to describe the conceptual bases of probabilistic reasoning and its generalizations, notably the Dempster-Shafer calculus of evidence, our attention during the reported research was primarily focused upon the formal characterization of possibilistic (i.e., "fuzzy logic") methods according to the similarity-based model that is briefly described above.

The major result of this research was a semantic model that was summarized in a number of publications and presentations [1,7,8,9,10] and that is discussed in detail in a technical

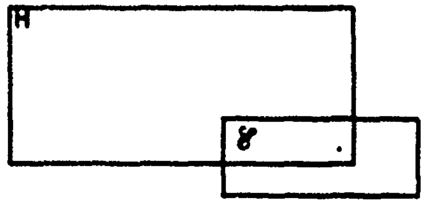

Figure 2: The approximate reasoning problem.

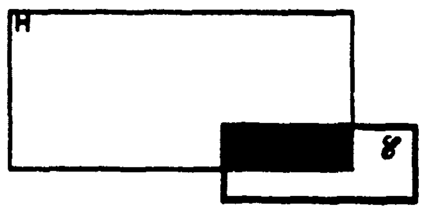

Figure 3: The probabilistic reasoning approach.

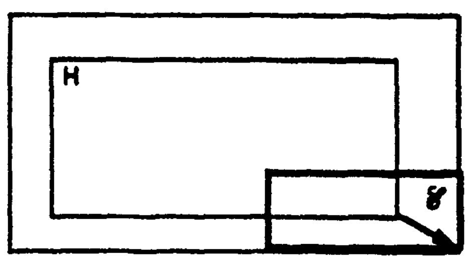

Figure 4: The possibilistic reasoning approach.

note [1], soon to be published in the *International Journal of Approximate Reasoning*.

The major characteristics of this model are:

- its ability to describe possibilistic techniques as the result of imposing metric structures upon a set of possible worlds rather than as the consequence of defining certain set measures (i.e., probabilities) on that set.
- its characterization of the metric properties of similarity or resemblance functions using operators previously considered only in the context of multivalued logics and the theory of probabilistic metric spaces (i.e., *triangular norms*).
- its description of metric relations between pairs of possible states or scenarios using well-known topological concepts (i.e., the *Hausdorff* distance), and the identification of relationships between such notions and the notions of unconditioned and conditional *possibility* distributions.
- the validation of the *generalized modus ponens*—the major inferential procedure of fuzzy logic—as a generalization of its classical counterpart,
- its ability to provide cogent descriptions of approximate relations between system variables.

Ongoing research, to be reported in the immediate future, is currently concerned with the following issues:

- *Derivation of similarity measures from possibility measures.* The semantic model described above has clearly established that possibilistic logic procedures rely on notions of similarity between plausible states of the world rather than on measures of the relative likelihood of such possibilities. While this model was developed primarily to improve understanding of fundamental conceptual matters, the relations that were uncovered during such development have significant implications of a practical nature. Of particular importance is the potential ability to derive similarity measures - -the bases for such analogical processes as case-based reasoning—from possibility distributions the formal expression of important qualitative physical laws. We have developed initial formulations for the derivation of such similarity measures on the basis of a formal result of L. Valverde on the representation of similarity measures.
- *The role of the notion of negation in possibilistic logic.* Conventional modal logics are concerned with the qualification of the truth of propositions by describing such truth as being either *necessary* (i.e., the unavoidable consequence of basic assumptions and the rules of logic), or *contingent* (i.e., the consequence of assumptions applicable to the particular situation under consideration). These considerations are the bases for the concepts of possibility and necessity, which related by a straightforward duality relation (based on the notion of negation) stating that something is possible if its negation is not necessary.

Our model of the semantics of fuzzy logic, while introducing graded (i.e., relative) notions of possibility and necessity based on measures of similarity, did not relate such notions using the concept of negation. Identification of such a relationship, however, is of significant conceptual and practical importance as knowledge of what is possible under certain circumstances may yield important information as to what is necessary in other cases.

Study of duality relations between pairs of subsets of possible worlds have led to the definition of new concepts of negation that are closely associated with the relations that exist between linguistic qualifiers that are *antonyms* of each other (e.g., [*rich*, *poor*] rather than [*rich*, *not-rich*]).

- *The study of the roles of system variables and concepts of independence in possibilistic logic.* We have studied the relations between similarity functions defined from the *joint* viewpoint of several variables (e.g., as when objects are differentiated using multiple attributes such as *color*, *volume*, *shape*) and *marginal* similarities that take only into account certain subsets of variables (e.g., measures of resemblance based solely on *color*). We have derived initial formulations for the derivation of joint similarity measures from their marginals and viceversa.

  Furthermore, study of the relationships that hold between similarity measures defined from diverse viewpoints have led to the definition of possibilistic measures of *independence* (or interaction) between variables. The results, which will be reported in a technical note that is currently under preparation, are of major practical importance to simplify complex processes of possibilistic inference (i.e., providing a possibilistic counterpart to the probabilistic methods of network decomposition). We are currently investigating representation formulas to derive marginal similarity functions without having to resort to transitive extension (i.e., chaining) of certain nontransitive relations. Availability of such formulas will greatly improve the efficiency of inferential processes.

- We are also investigating the conceptual relations between the important decision-theoretic notions of *utility*, *cost*, *desirability*, and *preference*. The central idea, based on concepts proposed by Rescher (N. Rescher, "Semantic foundations for the Logic of Preference," in N. Rescher, editor, *The Logic of Decision and Action*, Pittsburgh, 1967), is that such notions may be logically formalized by measures that quantify our preference to be in certain states of the world rather than others. Preliminary results indicate that a utility-based model will provide an even broader formal basis for possibilistic logic, while relating such preference measures with the metric structures of our basic model.

- We have developed a possibilistic formulation for the control of the navigation and for obstacle avoidance by autonomous vehicles that is being currently tested in the context provided by the SRI Autonomous Mobile Agent Platform.

4 Probabilistic Reasoning

We have also continued to investigate various issues of probabilistic reasoning, focusing upon questions of validity and generality of the Dempster-Shafer calculus of evidence.

We have given special attention to the discussion of recent concerns, raised within the technical community, about the conceptual soundness of this approach. Our contribution to this exchange, intended primarily to clarify various confusions and misconceptions, was summarized in a paper presented at the Third International Conference on the Management of Imprecision and Uncertainty by Expert Systems [5], which is expanded upon in an unpublished manuscript [6], currently under submission that is enclosed as part of this final report.

We have also continued our previous research on generalized probabilistic methods emphasizing the study of issues related to the treatment of conditional and dependent evidence. We have determined that , for reasonable definitions of conditional evidence distributions in the context of the DS calculus of evidence, these distributions are such that their combination with unconditioned evidence usually results (even for simple examples) in probability bounds that cannot be expressed within the Dempster-Shafer framework.

In connection, with these investigations we have derived a preliminary formulation of the problem of combination of conditional and unconditioned distributions as a linear program. In general, however, the solutions of such a problem will not obey the axioms of the calculus of evidence. Currently, we are focusing our attention upon three major questions:

- the determination of cases where the result of evidential conditioning is a belief function.
- the approximation of results not satisfying evidential axioms by belief functions that do.
- the development of a more general evidential calculus based on the notion of lower and upper probabilities.

Research Works Published or Presented under ARO Sponsorship

[1] E.H. Ruspini. *On the Semantics of Fuzzy Logic*. Technical Note No. 475, SRI International, December 1989. International Journal of Approximate Reasoning, forthcoming, 1991.

[2] E.H. Ruspini. *Approximate Reasoning: Past, Present, Future*. Invited paper. Information Sciences, forthcoming, 1990.

[3] E.H. Ruspini. *The Semantics of Belief Functions*. Presented at the 8th. International Congress of Cybernetics and Systems, New York, 1990 (Abstract published in Proceedings).

[4] E.H. Ruspini. *Similarity-Based Models for Possibilistic Logics*. Proceedings, Third International Conference on the Management of Imprecision and Uncertainty by Expert Systems, Paris, France, July 1990.

[5] J. Lowrance, E.H. Ruspini, and T. Strat. *On Evidential Reasoning*. Proceedings, Third International Conference on the Management of Imprecision and Uncertainty by Expert Systems, Paris, France, July 1990 .

[6] J. Lowrance, E.H. Ruspini, and T. Strat. *Understanding Evidential Reasoning*. In submission.

[7] E.H. Ruspini. *The Semantics of Approximate Reasoning*. Proceedings of the Fourth International Symposium on Knowledge and its Engineering, Barcelona, Spain, May 1990.

[8] E.H. Ruspini. *Possibility as Similarity: The Semantics of Fuzzy Logic*. Proceedings of the 1990 Conference on Uncertainty in Artificial Intelligence, Cambridge, Massachussets, 199C.

[9] E.H. Ruspini. *Similarity-based Interpretations of Fuzzy-Logic Concepts*. Proceedings of the 1990 International Conference on Fuzzy Logic and Neural Networks, Iizuka, Kyushu, Japan, 1990.

[10] E.H. Ruspini *The Semantics of Fuzzy Logic*. Proceedings 1990 NASA-JSC Symposium on Fuzzy Logic and Neural Networks, Houston, Texas, May 1990.

(To be published in Information Sciences, 1990)

APPROXIMATE REASONING: PAST, PRESENT, FUTURE

Technical Note No. 492

June 27, 1990

By: Enrique H. Ruspini
Artificial Intelligence Center
Computer and Engineering Sciences Division

This work was supported in part by the Air Force Office of Scientific Research under Contract No. F49620-89-K-0001 and in part by the United States Army Research Office under Contract No. DAAL03-89-K-0156.

Abstract

This note presents a personal view of the state of the art in the representation and manipulation of imprecise and uncertain information by automated processing systems. To contrast their objectives and characteristics with the sound deductive procedures of classical logic, methodologies developed for that purpose are usually described as relying on Approximate Reasoning.

Using a unified descriptive framework, we will argue that, far from being mere approximations of logically correct procedures, approximate reasoning methods are also sound techniques that describe the properties of a set of conceivable states of a real-world system. This framework, which is based on the logical notion of possible worlds, permits the description of the various approximate reasoning methods and techniques and simplifies their comparison. More importantly, our descriptive model facilitates the understanding of the fundamental conceptual characteristics of the major methodologies.

We examine first the development of approximate reasoning methods from early advances to the present state of the art, commenting also on the technical motivation for the introduction of certain controversial approaches.

Our unifying semantic model is then introduced to explain the formal concepts and structures of the major approximate reasoning methodologies: classical probability calculus, the Dempster-Shafer calculus of evidence, and fuzzy (possibilistic) logic. In particular, we discuss the basic conceptual differences between probabilistic and possibilistic approaches.

Finally, we take a critical look at the controversy about the need and utility for diverse methodologies, and assess requirements for future research and development.

Contents

1 Introduction

This note presents a personal view of the state of the art in approximate reasoning, the name used to describe several methodologies for the development of intelligent systems capable of manipulating imprecise and uncertain information.

Approximate reasoning techniques loosely based on the calculus of probability appeared almost simultaneously with the development of expert systems relying on classical (i.e., two-valued) logic techniques. Soon after these systems were introduced, other approaches to the treatment of uncertainty and imprecision were also proposed, both to generalize more or less conventional probabilistic schemes and to capture other aspects of imperfect knowledge, claimed to have a nonprobabilistic nature.

The short technological history of approximate reasoning methods may be described as being, from that moment, one of extreme controversy that has lasted to this day. Most of the proponents of classical probabilistic treatments, often described, although vaguely and somewhat misleadingly, as Bayesians,[1] have doubted the necessity for the introduction of other conceptual structures and have often sought to explain those frameworks in terms of probabilistic notions. Proponents of alternative approaches, on the other hand, have defended their techniques on the strength of two main arguments: the practical problems associated with the parameter-intensive procedures of conventional probability, often demanding knowledge of a large number of probability values; and, the nonprobabilistic nature of the uncertainties associated with the use of vague concepts.

Much of this disagreement has been clearly caused by misunderstandings about the fundamental philosophical characteristics of each approach. Lacking a suitable basis to interpret certain concepts, particularly those related to the "degrees of truth" of multivalued logics, it has been impossible, until recently, to provide an adequate framework to discuss fundamental issues in a rational manner.

This position paper on the past evolution of the field, its present state of the art, and desiderata for future evolution is the result of recent research by the author in basic semantic issues that are germane to the foundation of approximate reasoning. The presentation is based on the use of a central unifying framework: a formal model of the approximate reasoning problem that explains the similarities and differences between major methodologies. Using this "possible-worlds" model, we will also be able to compare the rationale of nonmonotonic logic approaches with that of approx-

[1]The qualifier Bayesian is used in the context of statistics to describe proponents of a statistical methodology and in the context of the philosophy of probability to denote various subjective views of probability. In Artificial Intelligence, the term has been loosely applied both to those investigating approaches based on the probability calculus and, more narrowly, to those espousing the decision-theoretic methods of subjective probability.

imate reasoning procedures. Although our model is a rigorous formalism, described in detail elsewhere [32,33] in connection with the logical foundations of the Dempster-Shafer calculus of evidence and fuzzy logic, our discussion will be kept as informal as possible to facilitate understanding our philosophical and technical position.

We will contend that regarding probabilistic and possibilistic approaches as competing alternatives is incorrect and confuses the need to describe different aspects of reality with the adequacy or ability of probability as a measure of likelihood. We will also take a critical look at the major claims supporting a narrow view of probability, based on a subjectivist interpretation that regards all forms of rational decision-making as necessarily demanding optimization of expected-utility functionals, and we dispute claims that only such approaches are endowed with either a suitable or a proven decision-theoretical apparatus.

On the basis of our theoretical arguments, and of recent success in the application of various techniques to practical problems, we will also argue that future accomplishment in the field lies in the rational development of tools leading to multiple complementary views of the implications of evidence rather than on arbitrary circumscription to a limited class of techniques and procedures.

2 The Development of Approximate Reasoning

Intelligent systems relying on approximate reasoning techniques [8,39] appeared in the 1970s, approximately at the same time as other systems seeking to emulate the expertise of specialists in diverse fields of endeavor. Problems related to the development of the expert systems based on classical deductive procedures, however, were primarily related to the need to organize knowledge and its processing in such a manner as to assure an efficient derivation of the truth value of hypotheses (i.e., either **true** or **false**). Systems such as MYCIN or PROSPECTOR— reasoning about medical and geological systems, where knowledge is limited and where observations may be difficult or impossible to make—were forced to deal, in addition, with issues that, to this day, have almost completely consumed the attention of approximate reasoning researchers.

These issues may be generally described as related to the extension of the basic derivation rule of classical logic, the modus ponens, which states that from the validity of an antecedent proposition p and that of the implication $p \rightarrow q$, it is possible to derive the validity of the consequent proposition q. Although a conventional expert system, using classical rules of derivation, could be assumed to have sufficient information to derive the validity of a hypothesis of interest, whenever knowledge was scarce or uncertain it was necessary to resort to other schemes that qualified in one way or another the meaning of the truth of propositions. Still imitating the network-oriented techniques of truth-value propagation of two-valued logic, the approximate reasoning schemes developed in early systems sought to propagate numeric truth values that were loosely related to probabilistic interpretations of uncertainty.

The concept of probability provides a most important tool to describe the state of systems that are known under less than desirable informational circumstances. Arising clearly from the need to make decisions despite undesirable knowledge handicaps, the notion of probability, seriously studied from the seventeenth century, has always played a major role in human judgment [16].

The appeal of probability as an instrument to assess system behavior is due to the empirically observed property that is expressed by the long-run stability of occurrence of certain events. Whether such a pattern of occurrence has been objectively quantified through experimentation or historical observation (objective interpretation), or is subjectively expressed by the willingness to gamble with certain stakes (subjective interpretation), it is clear that it provides a rational basis to formulate rational expectations about system state. Why would anybody, if such predictable stability of occurrence could not be assured, be willing to consciously bet on some outcomes rather than others if the real world defies any attempts to descriptive characterization?

Curiously enough, although probabilistic interpretations were always implicitly or explicitly intended by the developers of early approximate reasoning systems, and while the underlying calculi reflect such explanations, it seems also clear that the machinery of these devices was primarily oriented toward the emulation of the propagation schemes of classical logic with truth flowing from node to node through edges corresponding to implication rules. Approximate truth, measured by numbers associated with objective likelihood or expert confidence, also flowed from evidence to hypothesis in a scheme that generalized the true-false dichotomy of multivalued logic.

Regardless of the clearly intended probabilistic interpretations of those numbers, misgivings about their meaning and utility were sufficient to plant the seeds of the ensuing controversy. Concerns about the inability of probability to capture notions of evidential confirmation led the developers of MYCIN[39], for example, to introduce modified concepts ("certainty factors") as an alternative to direct use of conditional probabilities. In spite of subsequent studies showing that such certainty factors were related to probability values [18], it is clear that these worries were well founded, having been already eloquently expressed in the works of philosophers of science [34].

Although such concerns are indeed important and, despite some claims to the contrary, must still be properly addressed, other issues soon captured the attention of those seeking to develop expert systems with approximate reasoning capabilities. Beyond certain troublesome issues that were apparent when formulating the probabilistic calculi used by PROSPECTOR, arising from inconsistencies between "expert estimates" of probability values and the laws of probability, it was also clear to those engaged in the development of new expert systems that a typical application required estimation of a very large number of individual probability values [14], which were neither available or derivable from existing data.

In addition, other researchers, acquainted with the concepts and methods of multivalued logic [31,13], advanced the notion that some of the "degrees of truth" being propagated could be interpreted in a nonprobabilistic fashion. The theory of fuzzy sets, introduced by Zadeh in 1965 [45], had been for some time the focus of attention of these researchers and soon became a major source of techniques for the treatment of uncertainty by use of nonprobabilistic schemes.

The variety of approximate reasoning methods arising from this diversity—expressed as a preference toward either a variedly interpreted, more or less strict application of classical probability schemes; as approaches seeking the expression of ignorance about probability values, such as the Dempster-Shafer calculus of evidence; and as nonprobabilistic schemes like fuzzy logic— have led to a controversy that has endured to this day.

It has not been possible, until recently, to discuss these approaches with the help of a unifying framework that facilitates the interpretation of relevant concepts and the comparison of alternative methodologies. This unifying framework is based on a view of approximate reasoning problems as those wherein the truth-value of a hypothesis cannot be deduced from available information.[2] In other words, several scenarios, all consistent with evidence, may be conceived. In some of those stuations the hypothesis is true, while in others it is false.

The logical notion that we will use to characterize such conceivable states of affairs, situations, or scenarios, is the concept of " possible world" utilized by Carnap [4] in his logical treatment of the concept of probability, which was also employed by Nilsson [26] to derive a logic-based methodology for probabilistic reasoning.

3 Possible-World Models

A *possible world* may be briefly described as a function that assigns one and only one of the truth values **true** or **false** to every proposition (i.e., declarative statement) about the system that is being reasoned about. If we seek to describe and study the weather in Menlo Park, for example, the atmospheric conditions at several points in time are described by assigning specific values to meteorological variables such as temperature, humidity, and rainfall, or, equivalently, by assigning a truth value to

[2]Sometimes this characterization is extended to include those cases where that derivation is very difficult.

propositions such as

The temperature at 3PM was 75°F.

Since the value of system variables is unique (e.g., the temperature cannot be both 75°F and 85°F at the same time), it is clear that each possible world (i.e., an assignment of truth values) must satisfy certain consistency conditions that follow from the axioms of classical logic.

In approximate reasoning problems, however, we can usually do more to restrict the extent of the set of possible worlds that may conceivably describe the state of the system. Typically, the information or knowledge about the state of the system and its applicable rules of behavior, in spite of its defficiencies, is a major source of constraints that further limit the extent of the situations that must be considered. The subset of possible worlds that is logically consistent with this evidence is called the *evidential set*, and, in one form or another, is the concern of every approximate reasoning approach. In any approximate reasoning problem, by definition, some of these evidential worlds are such that a hypothesis is true in some of them and false on others, as depicted in Figure 1.

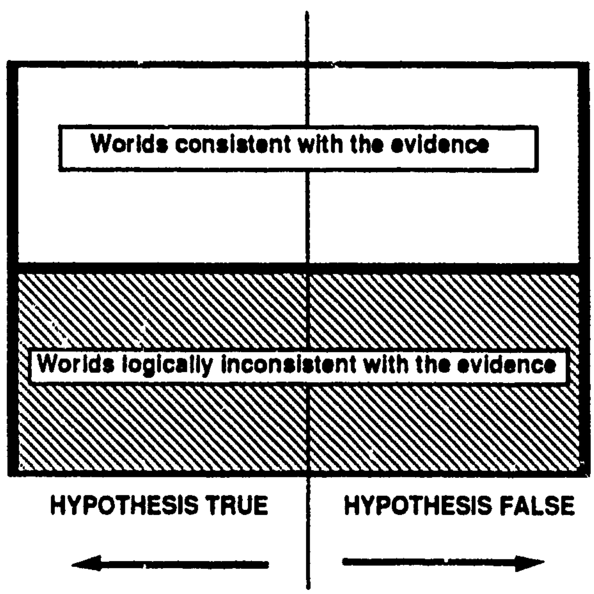

Figure 1: The approximate reasoning problem

The view of approximate reasoning problems that is afforded by this possible-world perspective also simplifies the understanding of the objective of approximate reasoning approaches. Lacking, by the nature of the problem, the ability to determine if the evidence implies whether we are in a situation where a hypothesis is true or in one where it is false, every approximate reasoning methodology seeks answers to a different problem: that of describing certain properties of the evidential set.

4 The Semantics of Approximate Reasoning

Our view of approximate reasoning methods as techniques to describe the evidential subset[3] e of possible worlds that are consistent with available information now allows a more detailed look into their philosophical bases.

Probabilistic methods, regardless of their subjective or objective semantics, seek to estimate measures of the subsets of the evidential set where a hypothesis h is true and where it is false, i.e., the values

$$\mu(h \wedge e) \quad \text{and} \quad \mu(\neg h \wedge e),$$

or other related quantities, such as likelihood ratios or conditional measures with respect to the evidential set ϵ. The measure μ is, however, an aggregate measure of set extension based on the additive law

$$\mu(p) + \mu(q) = \mu(p \wedge q) + \mu(p \vee q),$$

stating that its value over a set may be derived from knowledge of its value over a partition of nonintersecting subsets. Regardless of the mechanism used to derive the weights associated with individual members of the subsets, it should be clear that interactions and associations between possible worlds (e.g., distances) do not play any role in such quantities. Simply stated, all that matter are the weights of each individual point (more generally, each atomic subset) that are then added to gauge the extent of the subset.

Possibilistic methods, on the other hand, are based on notions of proximity and resemblance between pairs of possible worlds. This association or similarity is also a measure, albeit not one that may be expressed in terms of individual weights. Exploiting the idea that, in many systems, statements that are true in certain situations remain approximately true in similar instances (e.g., clothing that is appropriate when the temperature is 75°F will work nearly as well at 78°F), the purpose of possibilistic techniques is to describe the evidential set in terms of the similarity of its component possible worlds to other possible worlds used as reference landmarks.

The basic difference between probabilistic and possibilistic methods, therefore, goes beyond the use of different formulas to derive truth values. The methodologies are based on different conceptual approaches to the description of the evidential set;

[3]For simplicity, we refer loosely to sets and propositions are if they were the same objects.

they stress, in probabilistic reasoning, relative measures of set size, such as the ratio of previously observed true and false cases, while, in possibilistic reasoning, they stress binary measures of similarity that describe how far is any conceivable scenario from certain significant situations.

In both approaches, however, the objective is the description of properties of the evidential set rather than of any of its particular members. By contrast, certain nonmotonic logic techniques such as circumcription [24] rely on methods to choose least-exceptional worlds in the evidential set by extension of the "close-world assumption" [30], i.e., the only propositions or predicates that are true are those that are known to be true. These techniques may be considered general procedures to represent states of evidential knowledge by choice of prototypical situations. New evidence, however, may force retraction of some of the assumptions leading to the selection of other evidential worlds as prototypes. Another class of nonmonotonic reasoning techniques, while generally fitting the description given above, relies on prespecified "default" rules [29] to control the choice of prototypical worlds. Since these rules are usually formulated on the basis of plausibility notions rooted on statistical information (as in the famous example of Tweety and the flying ability of most live birds) it is not surprising that the derivation techniques and rules of these *preferential logics*—a name indicating their definition of a preferred order for models of a situation—resemble those of probabilistic reasoning. In fact, recent developments strongly point to the existence of a common unifying interpretation for both [28,15].

4.1 Probabilistic Reasoning

There can be little argument from any quarter that frequencies of occurrence of events satisfy the famous additive law that is axiomatized in the definition of set measure [17]. If propositions that describe event occurrence can only be assigned one and only one of the classical probability values, then it is obvious that whenever such repetitive occurrences are counted, then the sum of positive and negative occurrences must add up to the total number of relevant cases. As far as this objectivist interpretation of probability is concerned, therefore, there is little doubt that classical formalisms provide a suitable conceptual tool to capture the behavior of systems that expresses itself, as experimentally observed, in the form of stable frequency values.

Probabilities, viewed from the perspective of our possible-worlds model, may be considered as the basis of methods providing answers to a question that is related to but different from the undecidable issue of the validity of a hypothesis. Unable to state, because of lack of information, that h is either true or false, we describe instead the behavior of the system in the long run, by calculating the frequency of occurrence under similar circumstances.

Probabilistic reasoning schemes may be generally described as concerned with the computation of the joint probability distribution of several system variables, based on knowledge of the values of related marginal and conditional probability distributions. Whenever the required values are available it is possible, conceptually at least, to derive the required joint distributions. In fact, it may be fairly stated that, once it was understood that such derivation should be the goal of probabilistic reasoning systems, the attention of proponents of that methodological perspective has been almost completely directed toward the development of methods to simplify the required knowledge organization and manipulation [27].

Substantial concerns arise, however, regarding what must be done when the needed probability values are not known. In applied science, when unknown systems and phenomena are investigated, experiments are designed and performed to determine the basic laws of system behavior, which are typically expressed through quantitative relationships. If, based on such knowledge, rational courses of action are chosen, the careful scientist is then able to explain and justify his decisions on the basis of a strong epistemological apparatus supported both by empirical observation and by rational deduction. This scheme, which proceeds from information acquisition to decision making, embodies the experimental method of modern science. From such a perspective, probabilistic laws describe certain aspects of system behavior described by parameters that are estimated using the same methods that are universally accepted and employed in applied science.

Another view of probability, however, regards probability values as expressions of the degree of belief of rational decision makers regarding the validity of hypotheses. This degree of belief is quantified by the amount of money that a rational gambler is willing to bet in a gamble where the payoff, if the unknown truth value turns out to be **true**, is \$1. The probabilistic behavior of these degrees of belief is justified by a number of axiomatic systems [6,35] providing formal support not only to this subjectivist interpretation of probability but also to a decision-making methodology based on the maximization of expected utility. Related axiomatic formulations have been also developed to support the contention that the only correct procedure for updating such beliefs is the Bayes-Laplace rule [5]:

$$\mathbf{Prob}(q|p) = \frac{\mathbf{Prob}(p|q)\,\mathbf{Prob}(q)}{\mathbf{Prob}(p)}.$$

A number of researchers have questioned, in the past, the purportedly rational nature of these axiomatic systems. Their misgivings, which we share, arise both from questions about the rationality of some specific axioms, as noted by Suppes [42], and from observation of the behavior of rational decision-makers(including developers of the axiomatic formalisms) that contradicts the sure-thing principle, as observed by by Allais [1] and Ellsberg [11]. Kyburg [21] has also raised substantial concerns about the epistemological status and soundness of the subjectivist approach. The axiomatic system of Cox has also been criticized for its assumption that beliefs are measured by a single number [10] and, again, for the less-than-natural character of some axioms [38].

Proponents of this stringent orthodoxy have often argued that behavior departing from their theoretical requirements, however prevalent, is actually irrational. Such a claim, however, suffers from a fundamental methodological flaw. Rationality should be defined in terms of basic requirements that demand proper consideration of two fundamental factors: observed empirical evidence and the laws of logic. By requiring compliance with certain basic tenets of rational behavior, such as the famous avoidance of "dutch books," subjectivist schemes certainly attempt to meet one of these requirements, albeit in a limited fashion, as pointed out by Kyburg [21]. By defining rational behavior as that which results from utilization of the proponent's favorite scheme, the characterization of rationality is subjected to a curious argument that inverts the identity of what is rational with what must be done to ensure rational behavior. This inversion effectively ensures that the expected utility approach would always be considered to be rational: in fact, if any other behavior is observed, it would be, by definition, irrational.

This inversion of premises and conclusions is also apparent in other arguments, based on pragmatic necessity considerations, for the superiority of the subjectivist approach. If decisions, even those to obtain more information, must be made, then the elements required to make the decision (i.e., utility functions and degrees of belief) must be assessed. Conversely, any decision implies that such values have been, whether knowingly or not, chosen in some form or fashion. As a result of this close relation between the assessment of situations and the selection of suitable courses of action, guaranteed by the fact that values of expected utilities (i.e., numbers) may always be totally ordered, it is claimed that the subjectivist approach is the only one among approximate reasoning methods that has a rational decision-theoretic apparatus.

As appealing as such claims may be to some decision-makers, we must note again a curious exchange of roles in the scientific discovery process: decisions no longer follow from empirical observation and rational cogitation: rather, parameters that describe knowledge follow from a practical need to choose suitable actions. However pressing may be the need to derive decisions it should be clear that, in the absence of information, it is usually impossible to determine what is the best course of action. Any randomizing device would, under such circumstances, provide a total ordering of possible choices but there is very little to assure us that any behavior based on such arbitrary basis ought to be called rational.

The ultimate goal of an intelligent system is to take actions based on knowledge about the actual rather than the believed behavior of a real world system. It is difficult to see why, as noted by Kyburg [22], the latter should be given much attention outside psychological research. If applied science is, as generally admitted, a rational enterprise that seeks to uncover the secrets of the universe and to provide guidelines to take actions based on such knowledge, then it is clearly desirable that intelligent agents, in their quest for similar objectives, follow as closely as possible the essential procedures of the scientific method. The ability to produce decisions regardless of the extent and pertinence of available knowledge should be regarded as a handicap rather than as an advantage of a procedure: a fact readily noticed by those engaged in the solution of important real life problems [12]. As we pointed out before, whenever such knowledge is acquired, it is typically reported using a format that emphasizes the quality of the observational method and the strength of the arguments leading from empirical data to the author's conclusions rather than on the basis of personal confidence expressed by willingness to take gambling risks.

I have made a rather long exposition about the dichotomy between subjectivist and objectivist approaches to probability primarily because I believe this to be a major cause of a controversy that, beyond considerations that are solely germane to probabilistic reasoning, extends to the need for techniques that are not directly based on subjectivist orthodoxy. I have also been motivated by the desire to clearly expose a personal position that is shared by many in the approximate reasoning community but that is also often misleadingly described as being antiprobabilistic.

Far from being antagonistic to one approach for the simple sake of promoting others, my eclectic view is the direct result of practical experience with the development of models of complex systems, and of close familiarity with the application of mathematics to technological problems. Probability is indeed a powerful tool to describe chance-related aspects of the behavior of real-world systems. Recent contributions of probabilists and decision scientists, within and without the context of AI, such as the development of network-oriented procedures for probabilistic reasoning [27], are most important additions to our methodological arsenal.

There are, however, limitations on the capabilities of any tool, whether for system analysis or for any other purpose. As is true of any tool, including all methodologies described in this note, the applicability of probability is limited by its inability to perform functions that lie outside its scope, and by practical constraints on our ability to use it in specific situations. In spite of its unquestionable utility, other approaches also play a significant role in the description of the possible state of affairs. These techniques must not be considered to be competitors of probability but, rather, complementary techniques to enhance the understanding of the real world.

4.2 Generalized Probabilistic Reasoning

Those who worry about the potential lack of applicability of techniques based on conventional probability formalisms do not question the conceptual validity of probability as the appropriate tool to measure the frequency of occurrence of diverse events under various conditions or, in some cases, the strength of belief of decision-makers. Concerns about the problems caused by ignorance of probability values, however, have been expressed continuously since the nineteenth century by such prominent logicians as George Boole [3], and have led to the development of approaches to represent probabilistic ignorance by using subsets of possible probability values.

If, for example, the probability of validity of a proposition p is unknown, an interval probability method will represent such ignorance by assigning the interval $[0,1]$ as the value of the missing probability. If it is known, on the other hand, that an event has better than even chances of occurring, such knowledge will be represented by the $[0.5,1]$ interval. More generally, probabilistic knowledge may be represented as a set of possible probability values in a hyperdimensional cube, as in the convex probabilities approach of Kyburg [20].

The corresponding probabilistic calculi are straightforward conceptual extensions of the classic, number based calculus. Such extensions produce, for example, intervals of expected utility values on the basis of knowledge expressed as set of possible probability values. These intervals may be used, in many instances, to rank decisions in the same way that such choices are ordered with number-based schemes. When this ordering is not possible (e.g., overlapping intervals show that under certain scenarios A is preferrable to B, while, in other situations, B is to be preferred), the lack of a clear choice does not imply that the decision-theoretic apparatus is defective. Rather, the methodology is rich enough to tell us precisely how far empirical knowledge, combined with the laws of rational thought, can take us. If, beyond that point, it is imperative to do something—a rather unfortunate set of events—any selection scheme, from that point on, will be as rational as any other (i.e., very little).

Although the manipulation of intervals and sets of possible probability values alleviates some conceptual worries, it hardly helps in terms of the ability to perform the required computations. The situation, unfortunately, is made worse by the need to represent and manipulate probability bounds for subsets without the simplifying help that additivity provides for actual probability values. This unfortunate state of affairs is the primary reason for the popularity that an approach—capable of being interpreted in terms of interval probabilities— enjoys today as one of the major methodologies of approximate reasoning. This approach is the Dempster-Shafer calculus of evidence.

Originally developed by Dempster [7] in the context of statistical studies, the approach was further developed by Shafer [36] as a non-Bayesian alternative to the representation and manipulation of degrees of belief. Recently [32], application of possible-world semantic models to the interpretation of its major structures has shown that the approach is fully consistent with the classical calculus of probability, including the Bayes-Laplace formula. Smets [40] has also recently reviewed the structures of the calculus of evidence proposing, in addition, unconventional extensions based on a nonprobabilistic concept of belief.

The calculus of evidence may be readily understood using our basic model if it is recalled that, whenever assessing the validity of a hypothesis on the basis of emprical knowledge, there are three possible logical outcomes of any reasoning process: the hypothesis may be proved to be true, the hypothesis may be proved to be false, or the information may be insufficient to make either of those conclusions.

If the notation $\mathbf{K}p$ is used to denote the set of situations, i.e., possible worlds, where p can be proved true, if $\mathbf{K}\neg p$ correspondingly denotes those cases where it can be proved false, and if $\mathbf{I}p$ denotes the set of situations where the truth value of p cannot be established without ambiguity, then it is obvious that any probability function $\mathbf{Prob}(\cdot)$ will satisfy the equation

$$\mathbf{Prob}(\mathbf{K}p) + \mathbf{Prob}(\mathbf{K}\neg p) + \mathbf{Prob}(\mathbf{I}p) = 1 .$$

Furthermore, since the probability of $\mathbf{I}p$ may be positive, it will be true, in general, that

$$\mathbf{Prob}(\mathbf{K}p) + \mathbf{Prob}(\mathbf{K}\neg p) \leq 1 .$$

The calculus of evidence is based on the representation of the probabilistic information conveyed by evidence by means of *belief functions*. These functions may be readily interpreted in terms of the above probabilities of provability through the equation

$$\mathbf{Bel}(p) = \mathbf{Prob}(\mathbf{K}p) .$$

More importantly, these belief functions are usually expressible in a compact form by means of *basic probability assignments* or *mass functions*. These functions m, which are also defined over propositions, are related to belief functions by the equation

$$\mathbf{Bel}(p) = \sum_{q \Rightarrow p} m(q) .$$

The ability to represent and manipulate probability intervals by means of mass functions is the major reason for the appeal of the Dempster-Shafer methodology.

Although, in a typical decision problem, we are interested in the truth of p rather than its provability, lack of adequate information precludes determination of the probability of such truth. In general, however, it may be said that

$$\mathbf{Bel}(p) \leq \mathbf{Prob}(p) \leq 1 - \mathbf{Bel}(\neg p) .$$

Furthermore, these bounds cannot be improved.

This interpretation of the Dempster-Shafer calculus as concerned with probabilities of provability, as called by Pearl [27], was first formalized by the author using a possible-worlds model based on the use of a modal logic called epistemic logic. The formal system, which is equivalent to the modal system **S5** [19] used by Moore [25] in his pioneer work on the application of modal logic concepts to artificial intelligence problems, is enhanced by consideration of probability distributions over the set of possible worlds. In particular, the unary operator **K** represents the knowledge of a rational agent to prove that a proposition may be known or proved to be true.

The probability of the set of all possible worlds where a proposition p is the most specific proposition that is known to be true, called the *epistemic set*, corresponds to the values of the mass function. In any possible world, this most specific knowledge is the conjunction of all propositions that are known to be true in that possible world.

The semantic model of the Dempster-Shafer theory also validates the so-called Dempster's rule of combination, which permits the combination of belief and mass functions corresponding to evidential observations made under certain conditions of independence. When such conditions are not valid, use of this formula leads, of course, to erroneous results, often, although incorrectly, considered to be an essential handicap of the evidential reasoning approach, rather than a consequence of its misapplication.

From our perspective the only substantial example of such misapplication is that which results from improper use of the Dempster's rule of conditioning, i.e., a particular use of the rule of combination that is valid only under special circumstances, as a substitute for Bayes' rule. Certain methodological limitations of the calculus of evidence, notably the lack of methods to handle with sufficient generality the counterparts of conventional conditional probabilities, are more worrisome, in our opinion, than any distress arising from its misuse or its supposed lack of a decision-making apparatus.

4.3 Possibilistic Reasoning

Our basic semantic model also provides straightforward interpretations [33] for the major concepts and structures of possibility theory [46,9]: an approach to approximate reasoning derived from multivalued logics [31] and the theory of fuzzy sets [45]. The major formal tool that enhances our understanding of such structures is not a probabilistic measure of set size but, rather, a binary measure of proximity or distance, called a *similarity relation*.

Similarity considerations play a major role in human cognitive processes [44]. Informally, all such analogical processes are based on the notion that the validity of some propositions in a given situation extends also to other situations where the same basic conditions are prevalent.

In our model of possibilistic structures, the similarity between states of affairs is expressed by a function that assigns a number between 0 and 1 to every pair of possible worlds. The value of that function $S(w, w')$ for a pair of possible worlds quantifies the extent of resemblance between pairs of situations or scenarios, as evaluated from the viewpoint of the particular problem being considered. In a decision-making problem, for example, the decision maker may define such measures to describe the extent by which the consequences of certain decisions resemble desirable goals or objectives.

The highest similarity value, 1, indicates that, from the perspective of the system being studied, both situations are indistinguishable. The lowest value, 0, indicates that knowlege of what is true in one possible world does not help to derive what is true in the other.

Similarity scales are the measurement sticks used to describe the extent by which certain results may be extrapolated from one possible world to another. Unlike probability functions, which correspond to either measurable properties of physical systems or states of belief of rational agents, the similarity relations simply provide a mechanism to describe resemblance between states of affairs.

Similarity relations may also be regarded as generalizations of the modal-logic notion of *accessibility* or *conceivability* [19] by introduction of multiple binary relations R_α between possible worlds (one for each value of α between 0 and 1), defined by

$$R_\alpha(w, w') \text{ if and only if } S(w, w') \geq \alpha\,.$$

These relations also justify the use of a possibilistic terminology that regards propositions as being possible to some degree, thereby generalizing the classical definition of the modal operator for possible truth in a manner similar to that used by Lewis [23] in his treatment of counterfactual statements.

Certain requirements must be imposed to assure that similarity functions truly represent notions of resemblance between possible situations. Similarities between identical scenarios, for example, should have a value of 1, the highest possible value. Furthermore, if two different possible worlds are to be distinguished by means of similarity values, then it also makes sense to require that their similarity be strictly less than 1. It is likewise natural to require that the similarity between two particular scenarios be a symmetric function, i.e., w resembles w' as much as w' resembles w.

Beyond these properties of reflexivity and symmetry, it is also necessary to require that similarities satisfy a generalized form of transivity. If, given three possible worlds w, w' and w'', the worlds w and w' are highly similar while w' and w'' are also highly similar, it will be unreasonable to say that w and w'' may be highly dissimilar. The value of $S(w, w'')$ must, therefore, be bounded by below by a function of $S(w, w')$ and $S(w', w'')$, as expressed by the condition

$$S(w, w'') \geq S(w, w') \circledast S(w', w''),$$

which uses the binary operation $\circledast$ to denote the required function.

If certain reasonable requirements are imposed upon the function $\circledast$, it is easy to see that this function has the properties of *triangular norms*, which are usually introduced in multivalued logics [43] to relate the truth value of a conjunction $p \wedge q$ to the degrees of truth of p and q. These functions are motivated, in our model, by considerations that are related solely to metric concepts of proximity and resemblance. Important examples of triangular norms are given by the functions

$$a \circledast b = \min(a, b), \qquad a \circledast b = \max(a + b - 1, 0), \qquad \text{and} \qquad a \circledast b = ab,$$

called the *Zadeh*, *Lukasiëwicz*, and *product* triangular norms, respectively.

Similarity functions are trivially related by the relation

$$\delta = 1 - S,$$

to functions δ that have the properties of a distance or metric function. In the particular case where $\circledast$ is the triangular norm of Lukasïewicz, then δ is an ordinary metric or distance, which obeys the well-known triangular inequality

$$\delta(w, w'') \leq \delta(w, w') + \delta(w', w'').$$

If $\circledast$ is the Zadeh triangular norm, on the other hand, the transitivity property is equivalent to the stronger *ultrametric* inequality

$$\delta(w, w'') \leq \max(\delta(w, w'), \delta(w', w'')).$$

The structures introduced by similarity relations may be readily applied to generalize the subset inclusion relations that are the fundamental basis of deductive reasoning. These inclusion relations are typically expressed by conditional propositions of the form " If q, then p." stating that any state of affairs where q is true is such that p is also true. These conditional propositions, which permit the derivation of true propositions from knowledge of the truth of others by means of the rule of modus ponens, may be also stated using similarity structures by saying that any q-world has a p-world (i.e., itself) that is as similar as possible to it.

The ability to characterize proximity between possible worlds using a continuous scale of similarity provides for a more general characterization of the inclusion relations that hold between subsets of possible worlds (i.e., propositions). If the subset of q-worlds is not included in that of p-worlds, we may, however, use the similarity structure to quantify the amount of stretching required to reach a p-world from any q-world. The *degree of implication* function defined by the expression

$$\mathbf{I}(p \mid q) = \inf_{w' \vdash q} \sup_{w \vdash p} S(w, w'),$$

which is related to the well-known Hausdorff distance, provides such quantification as the size of the topological neighborhood of p that encloses q, as shown in Figure 2.

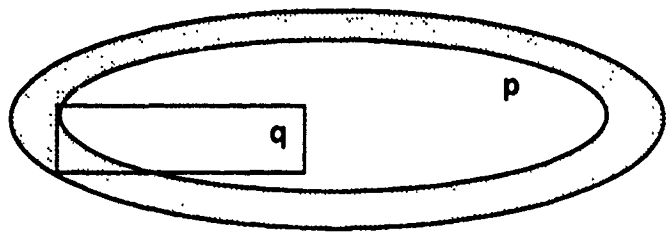

Figure 2: Degree of implication

The ability to express relationships between neighborhoods of different sets of possible worlds or, equivalently, between propositions permits the generalization of the modus ponens by use of the transitive property of the degree of implication function:

$$\mathbf{I}(p \mid r) \geq \mathbf{I}(p \mid q) \circledast \mathbf{I}(q \mid r),$$

illustrated in Figure 3.

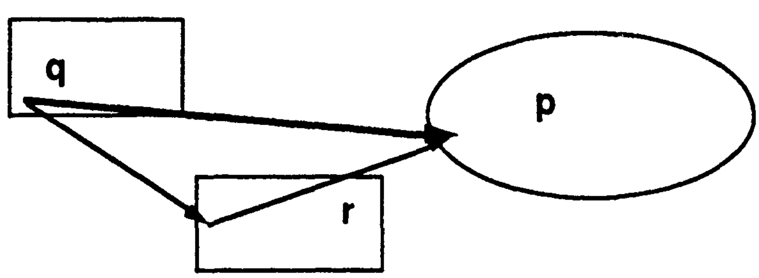

Figure 3: The generalized modus ponens.

The *generalized modus ponens* rule of Zadeh [46] is expressed by means of *possibility distributions*, which are themselves defined in terms of similarities between evidential worlds and those satisfying a given proposition p [33]. From the viewpoint of our similarity-based model, the generalized modus ponens may be thought of as a sound rule of logical extrapolation that exploits similarities between conceivable scenarios or situations. The fundamental topological structures that permit this type of reasoning are clearly different in character and nature than the measures of set extension that are the conceptual basis of probabilistic reasoning.

In closing, it is important to mention that posibilistic reasoning based on fuzzy logic has led recently to the implementation of a large number of successful commercial products [41]. These systems, which have primarily exploited the applicability of the technology to a variety of control devices, provide a clear indication of the usefulness of these ideas, which now also rest on clearly understandable theoretical foundations.

5 Looking ahead

The ability to explain the role and utility of the major approximate reasoning approaches by use of a unifying framework provides the rational basis to resolve most of the issues about relative importance and necessity. Rather than supporting any partisan contention about the superiority of one methodology over the others, this framework shows instead that a variety of tools are needed to produce effective descriptions of evidence and its implications.

Each methodology may play a significant role in every potential application of approximate reasoning techniques: a role that complements rather than substitutes forf other procedures. In the absence of compelling theoretical arguments for rejecting any approximate reasoning position and in the presence of substantial solid evidence of their usefulness and applicability, it is irrational to maintain positions that are needlesly divisive and polemic.

Recent investigations showing that there exist substantial functional rather than conceptual similarities between the network-oriented methods of conventional probabilistic schemes and the calculus of evidence [37], and indicating that fuzzy-set concepts and multivalued logic may be successfully blended to represent vague knowledge about probabilities [2], clearly point the way toward a more productive research collaboration between approximate reasoning specialists.

This collaboration should stress application of all valid concepts to the solution of practical problems rather than further continuation of the controversy about technological superiority or necessity. In particular, the example set by Japanese researchers in the development of a large number of commercial products of evident applicability illuminates the path that must be followed. The future lies in the solution of practical problems, both because of the direct importance of those problems, and because conceptual developments and clarifications usually follow, as is the case of the work discussed in this note, from the experiences gained producing such solutions. Having established needed conceptual bases to clarify controversial issues, we hope it is clear that this is the time to apply ideas rather than to continue to argue about them.

References

[1] M. Allais. Le Comportement de l'Homme Rationnel devant le risque: Critique des Postulats et Axiomes de l'Ecole Americaine. (*Behavior of a Rational Man under Risk: Critique of the Postulates and Axioms of the American School.*) Econometrica, 21:503–546, 1953.

[2] P. Bonissone and K. Decker. Selecting uncertainty calculi and granularity: an experiment in trading-off precision and uncertainty. In L.N. Kanaal and J.F. Lemmer, editors, *Uncertainty in Artificial Intelligence*, Amsterdam: North Holland, 1986.

[3] G. Boole. *The Laws of Thought*, 1854. New York: Reprinted by Dover Books, 1958.

[4] R. Carnap. *The Logical Foundations of Probability.* Chicago: University of Chicago Press, 1950.

[5] R.T. Cox. On inference and inquiry, an essay in inductive logic. In R.D. Levine and M. Tribus, editors, *The Maximum Entropy Formalism*, Cambridge, Massachussets: MIT Press, 1979.

[6] B. DeFinetti. La prévision: ses lois logiques, ses sources subjectives. *Annales de l'Institut H. Poincaré*, 7:1–68. 1937.

[7] A.P. Dempster. Upper and lower probabilities induced by a multivalued mapping. *Annals of Statistics*, 38:325–339, 1967.

[8] R.O. Duda, J. Gaschnig, J., P.E. Hart, K. Konolige, R. Reboh, P. Barrett, and J. Slocum. *Development of a computer-based consultant for mineral exploration.* SRI Project 5821 Final Report and SRI Project 6415 Annual Report. Artificial Intelligence Center, Menlo Park, California: SRI International, 1978.

[9] D. Dubois and H. Prade. An introduction to possibility theory and fuzzy logics. In P. Smets, E.H. Mamdani, D. Dubois, H. Prade, editors, *Non-standard Logics for Automated Reasoning*, 287–326, New York: Academic Press, 1988.

[10] D. Dubois and H. Prade. Comments on P. Garbolino, Probability: a Bayesian viewpoint. In P. Smets, M.R.B. Clarke, C. Froidevaux and J. Campbell, editors, *Positions and Debates in Nonstandard Logics*, New York: Academic Press, forthcoming.

[11] D. Ellsberg. Risk, ambiguity, and the Savage axioms, *The Quarterly Journal of Economics*, 75: 643–669. 1961.

[12] A.R. Feinstein. Clinical Biostatistics XXXIV: The haze of Bayes, the aerial palaces of decisions, and the computerized Ouija board. *Clinical Pharmacology and Therapeutics*, 21:482–496, 1977.

[13] B. Gaines. Fuzzy and probability uncertainty logics. *Inf. Control*, 38:154–169, 1978.

[14] T.D.Garvey, J.D. Lowrance and M.A. Fischler. An inference technique for integrating knowledge from disparate sources. *Proc. 7th. Intern. Joint Conf. on Artificial Intelligence*, Vancouver, British Columbia, 1981.

[15] H. Geffner. *Default Reasoning: Causal and Conditional Theories.* UCLA Cognitive Sciences Laboratory Technical Report (R-137), December 1989 (Ph.D. Dissertation, University of California, Los Angeles, 1989).

[16] I. Hacking. *The Emergence of Probability.* Cambridge, England: Cambridge University Press, 1975.

[17] P.R. Halmos. *Measure Theory.* New York: Springer-Verlag, 1974.

[18] D. Heckerman. Probabilistic interpretations for MYCIN certainty factors. In L.N. Kanaal and J.F. Lemmer, editors, *Uncertainty in Artificial Intelligence*, Amsterdam: North Holland, 1986.

[19] G.E. Hughes and M.J. Creswell. *An Introduction to Modal Logic.* New York: Methuen, 1972.

[20] H.E. Kyburg. *Logical Foundations of Statistical Inference.* Dordrecht: Reidel, 1974.

[21] H.E. Kyburg. Subjective probability: criticisms, reflections, and problems. *Journal of Philosophical Logic*, 7:157-180, 1978. Reprinted, with modifications, in *Epistemology and Inference*, Minneapolis: University of Minnesota Press, 1983.

[22] H.E. Kyburg. Interval valued probabilities. In P. Smets, M.R.B. Clarke, C. Froidevaux and J. Campbell, editors, *Positions and Debates in Nonstandard Logics,* New York: Academic Press, forthcoming.

[23] D. Lewis. *Counterfactuals.* Cambridge, Massachusetts: Harvard University Press, 1973.

[24] J. McCarthy. Circumscription: A non-monotonic inference rule. *Artificial Intelligence*, 13:27–40,1980.

[25] R. Moore. *Reasoning about Knowledge and Action.* Technical Note 191, Menlo Park, California: SRI International, 1980.

[26] N.J. Nilsson. Probabilistic logic. *Artificial Intelligence*, 28:71–87, 1987.

[27] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Mateo, California: Morgan Kaufmann, 1988.

[28] J. Pearl. Probabilistic semantics for nonmonotonic reasoning: a survey. *Proc. 1st. Intl. Conf. on Principles of Knowledge Representation and Reasoning (KR'89)*, Toronto, May 1989, 509–516, San Mateo, California: Morgan Kaufmann, 1989.

[29] R. Reiter. A logic for default reasoning. *Artificial Intelligence*, 13:81–132, 1980.

[30] R. Reiter. On close world data bases. In H. Gallaire and J. Minker, editors, *Logic and Data Bases*, 119–140, New York: Plenum, 1978.

[31] N. Rescher. *Many Valued Logic*. New York: McGraw-Hill, 1969.

[32] E.H. Ruspini. *The Logical Foundations of Evidential Reasoning*. Technical Note No. 408, Artificial Inteliigence Center, Menlo Park, California: SRI International, 1987.

[33] E.H. Ruspini. *On the Semantics of Fuzzy Logic*. Technical Note No. 475, Artificial Intelligence Center, Menlo Park, California: SRI International, Menlo Park,1989. *International Journal of Approximate Reasoning*, forthcoming.

[34] W.C. Salmon. Confirmation and relevance. induction, probability and confirmation. In G. Maxwell and R.M. Anderson, editors, *Minnesota Studies in the Philosophy of Science*, volume 6, Minnneapolis: University of Minnesota Press, 1975.

[35] L.J. Savage. *The Foundations of Statistics*. New York: Dover, second revised edition, 1972.

[36] G. Shafer. *A Mathematical Theory of Evidence*. Princeton, New Jersey: Princeton University Press, 1976.

[37] P.P. Shenoy and G. Shafer. An axiomatic framework for Bayesian and belief function propagation. *Proc. 4th. Workshop Uncertainty in Artificial Intelligence*, Minneapolis, Minnesota, 1988.

[38] G. Shafer. Comments on "An inquiry into computer understanding" by Peter Cheeseman. *Computational Intelligence*, 4:121–124, 1988.

[39] E.H. Shortliffe and B. Buchanan. A model of inexact reasoning in medicine. *Mathematical Biosciences*, 23: 351–379, 1975.

[40] P. Smets. Belief functions. In P. Smets, M.R.B. Clarke, C. Froidevaux and J. Campbell, editors, *Positions and Debates in Nonstandard Logics*, New York: Academic Press, forthcoming.

[41] M. Sugeno. *Industrial Applications of Fuzzy Control*. Amsterdam: North Holland, 1985.

[42] P. Suppes. The measurement of belief. *J. Roy. Stat. Soc. Ser. B.*, 36:160–175, 1974.

[43] E. Trillas and L. Valverde. On mode and implication in approximate reasoning. In M.M. Gupta, A. Kandel, W. Bandler, J.B. Kiszka, editors, *Approximate Reasoning and Expert Systems*, Amsterdam: North Holland, 157–166, 1985.

[44] A. Tversky. Features of similarity. *Psychological Review*, 84:433–460, 1977.

[45] L.A. Zadeh. Fuzzy sets. *Information and Control*, 8: 338–353, 1965.

[46] L.A. Zadeh. Fuzzy sets as a basis for a theory of possibility. *Fuzzy Sets and Systems*, 1:3–28, 1978.

(To be published in International Journal
of Approximate Reasoning, 1991)

ON THE SEMANTICS OF FUZZY LOGIC

Technical Note No. 475

December 15, 1989

By: Enrique H. Ruspini
Artificial Intelligence Center
Computer and Information Sciences Division

This work was supported in part by the Air Force Office of Scientific Research under Contract No. F49620-89-K-0001 and in part by the United States Army Research Office under Contract No. DAAL03-89-K-0156.

Abstract

This note presents a formal semantic characteri:ation of the major concepts and constructs of fuzzy logic in terms of notions of distance, closeness, and similarity between pairs of possible worlds. The formalism is a direct extension (by recognition of multiple degrees of accessibility, conceivability, or reachability) of the major modal logic concepts of possible and necessary truth.

Given a function that maps pairs of possible worlds into a number between 0 and 1, generalizing the conventional concept of an equivalence relation, the major constructs of fuzzy logic (i.e., conditioned and unconditional possibility distributions) are defined in terms of this generalized similarity relation using familiar concepts from the mathematical theory of metric spaces. This interpretation is different in nature and character from the typical, chance-oriented, meanings associated with probabilistic concepts, which are grounded on the mathematical notion of set measure. The similarity structure defines a topological notion of continuity in the space of possible worlds (and in that of its subsets, i.e., propositions) that allows a form of logical "extrapolation" between possible worlds.

This logical extrapolation operation corresponds to the major deductive rule of fuzzy logic —the compositional rule of inference or generalized modus ponens of Zadeh—an inferential operation that generalizes its classical counterpart by virtue of its ability to be utilized when propositions representing available evidence only match approximately the antecedents of conditional propositions. The relations between the similarity-based interpretation of the role of conditional possibility distributions and the approximate inferential procedures of Baldwin are also discussed.

A straightforward extension of the theory to the case where the similarity scale is symbolic rather than numeric is described. The problem of generating similarity functions from a given set of possibility distributions, with the latter interpreted as defining a number of (graded) discernibility relations and the former as the result of combining them into a joint measure of distinguishability between possible worlds, is briefly discussed.

Contents

List of Figures

List of Tables

To my friends Nadal Battle, Francesc Esteva, Ramón López de Mántaras, Enric Trillas, and Llorenç Valverde.

— En noblea són quatre coses especials e singulars.
Primerament que lo cavaller sia clar ens sos fets.
La segona que sia verdader.
La terça que sia fort de cor.
La quarta que haja coinexença,
car fort és odiosa descoinexença a Déu.

— *Joanot Martorell* [Martí Joan de Galba], *TIRANT LO BLANC* (CC).

1 INTRODUCTION

This note presents a semantic characterization of the major concepts and constructs of fuzzy logic in terms of notions of similarity, closeness, and proximity between possible states of a system that is being reasoned about. Informally, a "possible state" (to be formalized later using the notion of "possible world") is an assignment of a well-defined truth-value (i.e., either **true** or **false**) to all relevant declarative knowledge statements about that system.

The primary goal that guided the research leading to the results presented in this work has been one of conceptual clarification. A great deal of energy has been directed in past few years to debating the methodological necessity and relative merits of various approximate reasoning methodologies. As a result of these exchanges, the need to consider certain nonclassical approaches, has been questioned on a variety of bases.

Recognizing the need for the development of sound semantic formalisms that shed light on the nature of different approaches, the author has pursued, in the past few years, a line of theoretical research seeking to describe various approximate reasoning methodologies using a common framework. These investigations have recently shown the close connection between the Dempster-Shafer calculus of evidence [35] and epistemic logics. This relationship was elucidated by straightforward application of conventional probabilistic concepts to models of knowledge-states that distinguish between the truth of a proposition and knowledge (by rational agents) of that truth. Central to this development is the notion of "possible world" used by Carnap [6] to develop logical bases for probability theory.

The same central notion of possible state of affairs is also the conceptual basis of the results presented in this note, which is aimed at establishing the semantic bases of possibilistic logic with emphasis on the study of its possible relations and differences, if any, with probabilistic reasoning.

The results of this investigation clearly show that possibilistic logic can be interpreted in terms of nonprobabilistic concepts that are related to the notions of continuity and proximity. The major functional structures of fuzzy logic, i.e., possibility and necessity distributions,[1] may be defined in terms of the more primitive notion of similarity between possible states of a system using constructs that are the direct extension of well-known concepts in the theory of metric spaces. The topological metric structure that is so defined may be used to derive a sound inferential rule that is a form of logical "extrapolation." This rule is also shown to be the compositional rule of inference or generalized modus ponens proposed by Zadeh [53]. Conversely, possibility distributions—expressing resemblance from some specific regard—may be used to derive the actual similarity functions—discerning between possible worlds from the joint viewpoint of several respects.

The constructs that are used to derive the interpretation presented in this note are formally, structurally, and conceptually different from those that explain probabilistic reasoning, in either its objective or subjective interpretations, irrespective of methodological reliance on interval-based approaches to represent ignorance. The latter class of methods—measuring the relative proportion

[1] It is important to remark that the scope of this work is limited to the most fundamental concepts and constructs of fuzzy logic without examining related notions such as, for example, generalized quantifiers.

of (either observed or believed) occurrence of some event—are based on the mathematical notion of set measure, while the former—seeking to establish similarities between situations that may be used for analogical reasoning—are related to the theory of distances and metric spaces.

This presentation of the relationships between similarity-based concepts and possibilistic notions, while grounded on a formal treatment that is based on rigorous logical and mathematical formalisms, will be kept at a level that is as informal as possible. The purpose of this presentation style is to facilitate comprehension of major ideas without the clutter that would need to be otherwise introduced to keep matters strictly precise. For this reason, we will refrain from formal introduction of structures and axiom schemata, that, although correct and proper, may encumber understanding of the basic concepts.

Before we proceed to the detailed consideration of semantic models, I must briefly remark on the epistemological implication of these developments. The present interpretation is not claimed to be the only one that may be advanced to define the notion of possibility in terms of simpler concepts, nor do I claim that it may not be sometimes possible, even desirable, to model possibilistic structures from other bases. My intent is not to prove the conceptual superiority of one approach over another or to argue about the relative utility of different technologies. Rather, I hope that these results have contributed to establish the basic conceptual differences to the treatment of imprecise and uncertain information that are inherent in probabilistic and possibilistic methods; the former oriented toward quantifying believed or measured frequency of occurrence, and the latter seeking to determine propositions—implied by the evidence—that are similar, in some sense, to a hypothesis of interest. In other words, beyond accidental domain-specific relations, both types of methods are needed to analyze and clarify the significance of imprecise and uncertain information.

2 APPROXIMATE REASONING AND POSSIBLE WORLDS

Our point of departure is the model-theoretic formalisms of modal logics. Let us assume that declarative statements about the state, situation, or behavior of a real-world system under study are symbolically represented by the letters of some alphabet

$$\mathscr{A} = \{p, q, r, \ldots\},$$

which are combined in the customary way using the logical operators $\neg, \vee, \wedge, \rightarrow$ and $\leftrightarrow$ (to be interpreted with their usual meanings) to derive a language $\mathscr{L}$ (i.e., a collection of sentences). Furthermore, we augment this language by use of two unary operators N and Π, called the *necessity* and *possibility* operators, respectively, having usage governed by the rule

If ϕ is a sentence, then $N\phi$ and $\Pi\phi$ are also sentences,

introducing the ability to represent different modalities for the truth of propositions.

A *model* for this propositional system is a structure consisting of three components:

1. A nonempty set of possible worlds $\mathcal{U}$ introduced to represent states, situations, or behaviors of the system being modeled by our sentences. In what follows we will refer to this set as the *universe of discourse*, or *universe*, for short.

   We will also need to consider a nonempty subset $\mathscr{E}$ of the universe $\mathcal{U}$, which is introduced to model the set of conceivable worlds that are consistent with observed evidence. This set (possibly equal to the whole universe $\mathcal{U}$) will be called the *evidential set*. Throughout this note, we will assume that evidence about the world is always given by means of conventional propositions that allow to determine, without ambiguity, whether a possible world either is or is not a member of the evidential set.[2]

2. A function (called a *valuation*) that assigns one and only one of the truth values **true** or **false** to every possible world w in the universe $\mathcal{U}$ and every sentence ϕ in the language. Assignment of the truth-value **true** to a pair (w, ϕ) will be denoted $w \vdash \phi$ (i.e., ϕ is true in the world w).

   In what follows, we will use the same symbols to describe subsets of possible worlds and the propositions that are true only in worlds that are members of such subsets. For example, the symbol $\mathscr{E}$ will be used to denote both the evidential set and the proposition that asserts the validity of the corresponding evidential observations. Using this notation, for example, we will write $w \vdash \mathscr{E}$ to indicate that the world w is compatible (i.e., logically consistent) with the evidence $\mathscr{E}$.

   Furthermore, we will use the symbol $\mathscr{L}$, introduced above as a set of well-formed sentences, to denote also the power set of the universe $\mathcal{U}$. Rigorously, subsets of $\mathcal{U}$ strictly correspond to the classes of equivalence of the sentence set $\mathscr{L}$ that are obtained by equating logically equivalent sentences. In the same simplifying vein, we will drop also the customary distinction

[2] For the sake of simplicity, fuzzy evidential facts such as "Tom is rich," usually considered in fuzzy logic, will not be treated in this note. The meaning of such assertions will be discussed in a forthcoming paper.

between sentences—the linguistic expressions of something that may be true or false—and propositions—the actual things being asserted.

3. A binary relation R, between possible worlds, called the *accessibility*, *conceivability*, or *reachability* relation, introduced to model the semantic of the modal operators $\mathbf{N}$ and $\mathbf{\Pi}$.

It is not necessary to review here the well-known axioms [21] that restrict the assignment of truth values to well-formed sentences according to the rules of propositional logic. To facilitate comprehension of our formalism, we need to recall solely the rules that constrain assignment of truth values to sentences formed by prefixing other valid expressions with the modal operators, i.e.,

1. The sentence ϕ is *necessarily true* in the possible world w (i.e., $w \vdash \mathbf{N}\phi$) if and only if it is true in every world w' that is related to the world w by the relation R.

2. The sentence ϕ is *possibly true* in the possible world w (i.e., $w \vdash \mathbf{\Pi}\phi$) if and only if it is true in some world w' that is related to the world w by the relation R.

If, for example, the relation R relates worlds that share the same (possibly empty) subset of true sentences of the prespecified set of expressions

$$\mathcal{F} = \{\phi_1, \phi_2, \dots\},$$

i.e., $R(w, w')$ if and only if any sentence ϕ in $\mathcal{F}$ is either true in both w and w' or it is false in both w and w', then the resulting system has an "epistemic" interpretation that regards related possible worlds as "being possible for all we know" (i.e., observed evidence, corresponding to a subset of $\mathcal{F}$ is the same for both worlds). In this case, the necessity operator $\mathbf{N}$ corresponds the epistemic operator $\mathbf{K}$ of epistemic logics, with the corresponding system having the properties of the modal system S5, which was used—in the context of probability theory—as the semantic basis for the Dempster-Shafer calculus of evidence [35].

If, on the other hand, the original interpretation of logical necessity—corresponding to a relation R that is equal to $\mathcal{U} \times \mathcal{U}$, i.e., that relates every pair of possible worlds—is given to the operator $\mathbf{N}$, then a proposition is necessarily true if and only if it is true in every possible world.

If the relation R is chosen as

$$R = \mathcal{E} \times \mathcal{E},$$

then this interpretation may be used to characterize approximate reasoning problems as those where a hypothesis of interest is neither necessarily true nor necessarily false in worlds in the evidential set $\mathcal{E}$, reflecting the inability of conventional deductive techniques to unambiguously determine the truth-value of the hypothesis.[3]

In those problems, in spite of this fundamental impossibility, we may resort to approximate reasoning methods to describe various properties of the evidential set $\mathcal{E}$. For example, the probabilistic structures utilized by various probabilistic reasoning approaches typically characterize relations of the form

$$\mu(H \wedge \mathcal{E}) : \mu(\neg H \wedge \mathcal{E}),$$

between the "measures" of the subsets of the evidential set $\mathcal{E}$ where a hypothesis H is true or false, respectively.

[3]The notion of approximate reasoning problem is often extended to encompass situations where deductive techniques cannot always be used because of practical limitations on computational resources.

Our aim will be to study how other structures, defining a *metric* or *distance* in the universe $\mathcal{U}$, may be used to describe the nature of the evidential set. To do so, we will assign a different meaning to the accessibility relation, giving it an interpretation that regards related worlds as "similar" or "close" in some sense. We will require, however, a scheme that is richer than that provided by a single relation so that we can extend modal notions and derive semantics bases for fuzzy logic, which relies on concepts of degrees of matching or closeness expressed by real numbers between 0 and 1.

In what follows we will use the symbols $\Rightarrow$ and $\Leftrightarrow$ to denote strong implication and equivalence, respectively. A proposition q *strongly implies* p (denoted $q \Rightarrow p$) if and only if p is true in any world where q is. Similarly, p is *logically equivalent* to q (denoted $p \Leftrightarrow q$) if and only if p and q are true in the same subset of worlds of $\mathcal{U}$.

Following traditional terminology, we will say also that a proposition p is *satisfiable* if there exists a possible world p such that $w \vdash p$.

3 EXTENDED MODALITIES

We turn first our attention to the problem of generalizing modal logic formalisms to explain the structures and functions of fuzzy logic.

A number of authors have studied various relations between fuzzy and modal logics. Lakoff [24], Murai et al. [28], and Schocht [36] have proposed graded generalizations of basic modal constructs. Dubois and Prade [13,14] have also explored analogies between these nonstandard logics. In a recent paper [12], they have developed, in addition, a modal basis for possibility theory by means of the introduction of fuzzy structures into modal frameworks with the goal of deriving proof mechanisms that may be used in possibilistic reasoning.

The goal for the model presented in this note is somewhat different from the objectives guiding those efforts. We will seek explanations for possibilistic constructs on the basis of previously existing notions rather than generalizations of modal frameworks by means of fuzzy constructs. The model presented here is not based on the use of graded notions of possibility and necessity as primitive —and, by implication, easy to understand—structures. The foundation for this model is provided by a generalization of the accessibility relation, which is given a simple interpretation as a measure of resemblance and proximity between possible worlds.

We will extend the notion of accessibility relation to encompass a family of nonempty binary relations R_α that are indexed by a numerical parameter α between 0 and 1. These relations, which are nested, i.e.,

$$R_\alpha \subseteq R_\beta, \quad \text{whenever } \beta \leq \alpha,$$

are introduced to represent different degrees of similarity, using a scheme that is akin to that used by Lewis in his study of counterfactuals [25]. The family of accessibility relations introduced here differs from that proposed by Lewis, however, in its use of numerical indexes[4] and in the nature of the overall modeling goals that, in Lewis' formalism, are intended to represent changes of scale induced by consideration of different restrictive statements.

3.1 Similarity Relations

To facilitate the definition of a family of accessibility relations we introduce a *similarity function*

$$S : \mathcal{U} \times \mathcal{U} \longmapsto [0,1],$$

assigning to each pair of possible worlds (w, w') a unique *degree of similarity* between 0 (corresponding to maximum dissimilarity) to 1 (corresponding to maximum similarity).

With the help of this function, we will then say that w and w' are related to the degree α, denoted $R_\alpha(w, w')$, if and only if $S(w, w') \geq \alpha$. In this way, the relations R_α have the required nesting property with R_0 corresponding to the whole Cartesian product $\mathcal{U} \times \mathcal{U}$ (or, every possible world is at least similar in a degree zero to every other possible world).

[4] We will later see that similarities may be measured using more general, nonnumeric, scales. For simplicity reasons, we will avoid at this point the introduction of more general schemes that unnecessarily complicate the exposition.

Some properties are required to assure that the function S has the required semantics of a metric relationship capturing the intuitive notion of similarity or "proximity." It is first necessary to demand that the degree of similarity between any world and itself be as high as possible, i.e.,

$$S(w, w) = 1, \quad \text{for all } w \text{ in } \mathcal{U}.$$

This property assures that every one of the accessibility relations R_α will be *reflexive* and, following the nomenclature introduced by Zadeh for fuzzy relations [52], we will also say that the similarity relation is reflexive.

Next, we will call for the function S to be *symmetric*, i.e.,

$$S(w, w') = S(w', w), \quad \text{for any worlds } w \text{ and } w' \text{ in } \mathcal{U}.$$

This is a very natural requirement of any relation intended to represent a relation of resemblance between objects.

Finally, and most importantly, we will impose a form of *transitivity* requirement upon the similarity function S that turns it into a generalized equivalence relation. The purpose of this restriction is to assure that S has a reasonable behavior as a metric in the universe of possible worlds. It would certainly be surprising if, for some similarity S, we were to be told that w and w' are very similar and that w' and w'' are also very similar, but that w does not resemble w'' at all. Clearly, there should be a lower bound on the possible values of $S(w, w'')$ that may be expressed as a function of the values of $S(w, w')$ and $S(w', w'')$. We will express such a constraint using a numeric operation, denoted $\circledast$, that takes as arguments two real numbers between 0 and 1 and that returns another number in the same range, i.e.,

$$\circledast \colon [0,1] \times [0,1] \longmapsto [0,1],$$

in the form of the inequality

$$S(w, w'') \geq S(w, w') \circledast S(w', w''),$$

assumed valid for any worlds w, w' and w'' in the universe $\mathcal{U}$. Recurring again to a modal terminology, the above transitivity constraint, which will be called *$\circledast$-transitivity*, may be rewritten in relational form as

$$R_{\alpha \circledast \beta} \subseteq R_\alpha \circ R_\beta, \quad \text{for all } 0 \leq \alpha, \beta \leq 1,$$

making obvious its generalization of the conventional definition of transitivity for ordinary binary relations, i.e.,

$$R \subseteq R \circ R.$$

Since the role of $\circledast$, through recursive application, is that of providing a lower bound for the similarity between the two end members w_1 and w_n of a chain of possible worlds $[w_1, w_2, \ldots, w_n]$, it is obvious that the operation $\circledast$ should be commutative and associative. Furthermore, it should also be nondecreasing in each argument, as it is reasonable to ask that the desired lower bound be a monotonic function of its arguments. Finally, it is also desirable to ask that

$$\alpha \circledast 1 = 1 \circledast \alpha = \alpha,$$

i.e., that the values of the similarities of two indistinguishable objects to a third should be the same. These requirements are equivalent to demanding that the operation $\circledast$ be a *triangular norm* [37], or *T-norm*, for short.

Triangular norms, originally introduced in the theory of probabilistic metric spaces to treat certain statistical problems, play a distinguished role in $[0,1]$-multivalued logics [1,11,17,31] as the result of imposing reasonable requirements upon operations that produce the truth value of the conjunction of two expressions as a function of the truth values of the conjuncts. Furthermore, generalized similarity relations (called B-R relations by Zadeh [54]) also have an important function, to be examined further later in this note, in the generalization of the inferential rule of *modus ponens* [43,10]. Our axiomatic derivation for the requirement that $\circledast$ be a T-norm is based, however, solely on metric considerations, applied here to a space of possible worlds, but is valid in general metric spaces.

From the axioms of triangular norms, it is easy to see that

$$\alpha \circledast \beta \leq \min(\alpha, \beta),$$

showing that the minimum function, itself a T-norm, is the largest element in this class of operations. Its minimal element, on the other hand, is the noncontinuous function $\circledast$ defined by

$$\alpha \circledast \beta = \begin{cases} \alpha, & \text{if } \beta = 1, \\ \beta, & \text{if } \alpha = 1, \\ 0, & \text{otherwise.} \end{cases}$$

Every symmetric and reflexive relation is $\circledast$-transitive for this triangular norm, which is, therefore, of little practical utility.

In what follows, we will also impose a most reasonable additional assumption of continuity of $\circledast$ with respect to its arguments (i.e., why should there be a jump in the value of a lower bound provided by $\circledast$ when the values of its arguments are slightly changed?). The class of continuous T-norms does not have a minimal element, although under certain additional assumptions (requiring T-norms to be also *J-copulas* [37]), the inequality

$$\max(\alpha + \beta - 1, 0) \leq \alpha \circledast \beta$$

also holds true, showing that certain important continuous T-norms lie between that of the $\aleph_1$-logic of Lukasïewicz [17] and that of the original fuzzy logic proposed by Zadeh [53].

Continuous triangular norms play a significant part in the theories of pattern recognition and automatic classification. The author [33] proposed the use of generalized similarity relations based on the T-norm of Lukasïewicz to generalize existing classification techniques—based on the mapping of a similarity function into a conventional equivalence relation—to the fuzzy domain—by mapping these T-norms (called *likeness relations* by Ruspini) into generalized fuzzy partitions. Bezdek and Harris [3] independently studied axiomatic approaches to cluster analysis based on the use of several continuous T-norms.

The author has also studied [34] the possible relation between the multivalued logic and similarity related aspects of T-norms, and suggested that the degrees of similarity between two objects A and B may be regarded as the "degree of truth" of the vague proposition

"A is similar to B."

Having argued that S should have the structure of a generalized equivalence relation, we will assume, mainly for reasons of simplicity, that the function S is the dual of a "true" distance, i.e., that

$$S(w, w') = 1 \text{ if and only if } w = w'.$$

This restriction, which is not substantial, is introduced primarily to assure that different possible worlds may be distinguished by means of the function S. Otherwise, the equivalence relation that relates two worlds w and w' if and only if $S(w, w') = 1$ may be used to partition our universe $\mathcal{U}$ into "indistinguishable" nonintersecting classes—indicating that our metric cannot discriminate between significant differences in system state.

Before closing our presentation of generalized similarity relations, it is important to remark upon the close relation between the notion of similarity and that of distance. If a function δ is defined in terms of a similarity function S by the simple relation

$$\delta = 1 - S,$$

then it is easy to see that the function δ has the properties of a metric or distance. This is evident if the operation $\circledast$ corresponds to the T-norm of Lukasïewicz, since the transitivity condition is equivalent to the well-known triangular inequality, i.e.,

$$\delta(w, w'') \leq \delta(w, w') + \delta(w', w'').$$

If other T-norms are used, even stronger inequalities hold, with the so-called "ultrametric inequality"

$$\delta(w, w'') \leq \max\left(\delta(w, w'), \delta(w', w'')\right)$$

being valid for the T-norm of Zadeh. In this case, each of the relations in the family R_α (known in fuzzy set theory as the α-cut[5] of the similarity S) is a conventional equivalence relation. This fact was exploited, prior to the introduction of fuzzy set theory and fuzzy cluster analysis, by a variety of clustering procedures of the "single-link" type [22,40].

3.2 Possible and Necessary Similarity

Our semantic formalization needs require the introduction of constructs to indicate the extent by which a concept exemplifies, illustrates, or is an adequate model of another concept. Our interpretations shall, therefore, be oriented toward characterization of the degree by which a concept can be said to be a good example of another concept with the purpose of defining vague concepts by means of measures of proximity between defined and defining concepts. In our treatment, each of the multiple "definiens" will be a conventional proposition corresponding to a subset of possible worlds. It is conceivable, however, that new vague concepts might also be described by indicating their metric relations to other vague concepts.

The required constructs are based on the idea that whenever p and q are propositions such that $p \Rightarrow q$, then any p-world is an "example" of a q-world. This basic notion will be generalized by the introduction of modal structures that define to what degree possible worlds that satisfy a certain proposition q fit a vague concept. Some of those possible worlds are "paradigmatic" of the vague concept, i.e., they fit it to a degree equal to 1 in the same sense that we may say, for example, in an absolute (i.e., nongraded) sense that somebody whose height is 7 ft is definitely "tall." If we use a notion of graded fitness, however, certain worlds will fit the concept to a degree, i.e., they resemble (or are similar) to some paradigmatic example of the vague concept.

The conventional interpretation of possibility needs to be modified, therefore, to capture the idea that a particular possible world is similar in some degree to another world that satisfies a "reference" proposition.

[5]The α-cut of a fuzzy set μ: $\mathcal{U} \mapsto [0,1]$ is the conventional set of all points w such that $\mu(w) \geq \alpha$. A similar concept is defined for relations as subsets of a product space $\mathcal{U} \times \mathcal{V}$.

More generally, however, we will be interested in relations of similarity between *pairs of subsets* of possible worlds rather than between pairs of possible worlds. This requirement complicates matters considerably since we will be forced to consider both the "validity" of a proposition p in *some* world where another proposition q is true, as well as its applicability in *every* world where q is true. In the former case, we will care about the existence of q-worlds that are similar to some degree to some p-world, while in the latter we will be concerned with the size of the minimum neighborhood of p (as a subset of the universe $\mathcal{U}$) that fully encloses the subset q.

This dual concern for what may possibly apply and what must necessarily hold—an essential aspect of modal logic—is typical of situations where relationships between ensembles of objects are described in terms of relations between their members. In the probability calculus, for example, knowledge of probabilities over certain families of subsets provides "sharp" upper and lower bounds (called *inner* and *upper probabilities*, respectively) for the probabilities of other subsets—an important fact in the extension of set measures to larger domains [19]. The role and properties of these bounds in the Dempster-Shafer calculus of evidence is well-known, having been described in the original paper of Dempster [8], related to concepts of modal logic by Ruspini [35], and being also the subjects of considerable formal study [7] as mathematical structures.

Analogies between the role of probabilistic bounds (i.e., bounds for probability values) and possibility/necessity distributions—shown below to have play a similar part with respect to metric structures—have been the source of much of the confusion about the need for possibilistic schemes. Each upper/lower-bound pair, however, leads to a substantially description of the nature of a subset of possible worlds, being, in either case, measures that arise naturally when pointwise properties are extended to set partitions. General properties of these measures have been studied by Dubois and Prade [11] in the context of approximate reasoning and in other regards by Pavlak [30].

Our generalizations of the notions of possibility and necessity are related to the so-called *de re* [21] interpretation of the statement "If q, then p is possible" as the modal propositional relation

$$q \Rightarrow \mathbf{\Pi} p\,.$$

We will say that the proposition q *implies*, or is a *necessary model* of, the proposition p *to the degree* α if and only for every q-world w there exists a p-world w' that is at least α-similar to it, (i.e., $S(w, w') \geq \alpha$), or equivalently, whenever

$$q \Rightarrow \mathbf{\Pi}_\alpha p\,.$$

Similarly, we will say that the proposition q *is consistent with*, or is a *possible model* of, the proposition p *to the degree* α [6] if and only there exist a q-world w and a p-world w' that are at least α-similar, or equivalently, whenever

$$\neg(p \Rightarrow \neg\mathbf{\Pi}_\alpha q)\,.$$

The similarity function that we have introduced in the universe $\mathcal{U}$ provides us with a simple mechanism to quantify both the extent of "inclusion" and that of the "intersection" between pairs of subsets of possible worlds.[7]

[6] Note that our characterizations of both possibility and necessity distributions are based in the modal possibility operators $\mathbf{\Pi}_\alpha$.

[7] For reasons that by now should be evident, we will not need to introduce a concept of "unconditioned possibility" although it would be easy to do so using $q = \mathcal{U}$. Being concerned with the power of certain propositions to exemplify other conditions, we will not have much occassion to deal with the strength of tautologies in that regard.

3.3 Possibilistic Implication and Consistence

The notion of subset inclusion and its related concept of set identity are of central importance in deductive logic, since subsets of possible worlds are formally equivalent to propositions with subset inclusion and identity corresponding to logical implication and equivalence, respectively. These propositional relationships are the basis of derivation rules such as the modus ponens. The notion of intersection plays a similar role in modal analyses because of its ability to express the potential validity of a statement.

Classical accounts, however, recognize only two "degrees" of inclusion corresponding to the cases when either a set q is a subset of another set p or it is not, with a similar dichotomy applying to degrees of intersection. Our generalization exploits the metric structures defined between sets of possible worlds by introducing measures that describe a subset as enclosed in a *neighborhood* (of some size) of another set while intersecting another of its neighborhoods (of "smaller" size).[8] The problem of measuring the "size" of those neighborhoods is the subject of our immediate considerations.

3.3.1 Degree of Implication

Our definition of partial implication between propositions was based on conditions that determine whether, given two propositions p and q, one of them implies the other to the some value α. In particular, since every world w is always similar in a degree that is at least equal to zero to any other world w', it is always true that any proposition q implies any other proposition p to the degree zero. It is often the case, however, that the degree of implication between p and q is at least equal to some certain positive value α.

If we want to generalize procedures based on inclusion relationships, such as the modus ponens, in an efficient fashion, we will need measure the "optimal" (or maximum) value of the parameter α such that q implies p to the degree α. This value is a measure of the degree by which the set of all p-worlds must be "stretched" to encompass the set of all q-worlds. The least upper bound of the values of the similarities between any q-world w' and some p-world w (depending, in general, from w') is given by the *degree of implication* function:

Definition: The *degree of implication* of p by q is the value

$$\mathbf{I}(p\,|\,q) = \inf_{w'\vdash q}\ \sup_{w\vdash p} S(w,w')\,.$$

Defined in this way, the degree of implication $\mathbf{I}(p\,|\,q)$ is a measure of the "minimal amount" of stretching required to reach a p-world from any q-world, in the sense that if $\beta < \mathbf{I}(p\,|\,q)$, then

$$q \Rightarrow \Pi_\beta p\,.$$

Furthermore, α is the largest real value for which the above statement may be made.

As the following theorem makes clearer, this function provides the bases for the generalization of the modus ponens. This truth-derivation procedure may be thought of as an expression of the nesting relationships that hold between the sizes of neighborhoods of such subsets.

[8] It is important to recall that, due to our reliance on similarity rather than on the dual notion of dissimilarity or distance, high values of α correspond to low values of "stretching" or to smaller set neighborhoods.

Theorem: The degree of implication function,

$$\mathbf{I}: \mathcal{L} \times \mathcal{L} \mapsto [0,1],$$

has the following properties:

(i) If $p \Rightarrow r$, then $\mathbf{I}(p \mid q) \leq \mathbf{I}(r \mid q)$

(ii) If $q \Rightarrow r$, then $\mathbf{I}(p \mid q) \geq \mathbf{I}(p \mid r)$

(iii) $\mathbf{I}(p \mid q) \geq \mathbf{I}(p \mid r) \circledast \mathbf{I}(r \mid q)$

where p, q and r are any satisfiable propositions.

Proof: The first two properties are an immediate consequence of the definition of degree of implication. To prove the third, observe that by definition of similarity

$$S(w, w') \geq S(w, w'') \circledast S(w'', w')$$

for any worlds w, w', and w''.

Taking the supremum on both sides of this inequality with respect to all worlds $w \vdash p$, it follows, because $\circledast$ is continuous, that

$$\sup_{w \vdash p} S(w, w') \geq \left[\sup_{w \vdash p} S(w, w'')\right] \circledast S(w'', w').$$

Since this expression is true, in particular, for all worlds $w'' \vdash r$, it is true that

$$\begin{aligned} \sup_{w \vdash p} S(w, w') &\geq \left[\inf_{w'' \vdash r} \sup_{w \vdash p} S(w, w'')\right] \circledast S(\hat{w}, w') \\ &= \mathbf{I}(p \mid r) \circledast S(\hat{w}, w'), \end{aligned}$$

where $\hat{w}$ is any world such that $\hat{w} \vdash r$.

From this inequality, it follows, since $\circledast$ is continuous, that

$$\sup_{w \vdash p} S(w, w') \geq \mathbf{I}(p \mid r) \circledast \left[\sup_{\hat{w} \vdash r} S(\hat{w}, w')\right].$$

Taking now the infimum on both sides of this expression over all worlds w' such that $w' \vdash q$, it is easy to see, using again the continuity of $\circledast$, that

$$\inf_{w' \vdash q} \sup_{w \vdash p} S(w, w') \geq \mathbf{I}(p \mid r) \circledast \left[\inf_{w' \vdash q} \sup_{\hat{w} \vdash r} S(\hat{w}, w')\right],$$

proving the $\circledast$-transitivity of $\mathbf{I}$. ∎

Note, that since $\mathbf{I}(q \mid q) = 1$ for any proposition q, the following statement is also true:

Corollary. If p and q are propositions in $\mathcal{L}$, then

$$\mathbf{I}(p \mid q) = \sup_{r} \left[\mathbf{I}(p \mid r) \circledast \mathbf{I}(r \mid q)\right].$$

Notice also that if $\mathbf{I}(p \mid q) = 1$, then

$$\sup_{w \vdash p} S(w, w') = 1\,, \qquad \text{for all } w' \vdash q\,.$$

Under minimal assumptions (assuring that the supremum operation is actually a maximization), this relation is equivalent to stating that q strongly implies p, or that any q-world is also a p-world.

The nonsymmetric function $\mathbf{I}$ measures the extent by which every world w' in a certain class resembles some world w (dependent of w') in a reference class, possibly explicating the nature of the nonsymmetric assessments [45] found in psychological experimentation when subjects are asked to evaluate the degree by which an object "resembles" another. The results obtained in those experiments suggest that human beings, when assessing similarity between objects, use one of them (or a class of similar objects) as a reference landmark to describe the other. Such assymmetries might be explained by noticing that, in general, $\mathbf{I}(p \mid q) \neq \mathbf{I}(q \mid p)$, indicating that the stronger stimulus might generally be used to construct a reference class, which is then used to describe other stimuli.

The degree of implication of one proposition by another can be readily used to generate a measure of similarity between propositions that generalizes our original measure of similarity between possible worlds:

$$\widehat{S}(p, q) = \min\left[\mathbf{I}(p \mid q), \mathbf{I}(q \mid p)\right],$$

quantifying the degree by which the propositions p and q are equivalent.

It may be readily proved [44], from its definition and from the transitivity property of $\mathbf{I}$ that $\widehat{S}$ is a reflexive, symmetric, and $\circledast$-transitive function between subsets of possible worlds. This similarity function is the dual of the well-known *Hausdorff distance*, defined between subsets of a metric as a function of the distance between pairs of their members [9], which is given by the expression

$$\hat{\delta}(A, B) = \max\left[\left(\sup_{x \in A}\ \inf_{y \in B} \delta(x, y)\right), \left(\sup_{x \in B}\ \inf_{y \in A} \delta(x, y)\right)\right].$$

The result expressed by the transitive property of the degree of implication may be stated using modal notation in the form

$$q \Rightarrow \Pi_\alpha r \ \text{ and } \ r \Rightarrow \Pi_\beta q \quad \text{imply that} \quad q \Rightarrow \Pi_{\alpha \circledast \beta}\, p,$$

as the simplest form of the generalized modus ponens rule of Zadeh.

The relationship between this rule and the classical modus ponens is easier to perceive if it is remembered that classical conditional propositions of the form "If q, then p," simply state that the set of q-worlds is a subset of the set of p-worlds. Such relationships of inclusion may also be described in metric terms by saying that every q-world has a p-world (i.e., itself) that is as similar as possible to it.

Logic structures, however, only allow us to say that either q implies p or that q implies its negation $\neg p$, or that neither of those statements is true. By contrast, similarity relations allow measurement of the amount by which a set must be "stretched" (as illustrated in Figure 1) to enclose another set. Using such metrics, we may describe the generalized modus ponens as a relation between the stretching required to reach p from any point of the set r, the stretching required to reach r from any point of the set q, and the stretching required to reach p from any point of the set q.

In Section 5 we will derive alternative expressions for the generalized modus ponens that allow to propagate both measures characterizing degree of implication and *degree of consistence*; a dual

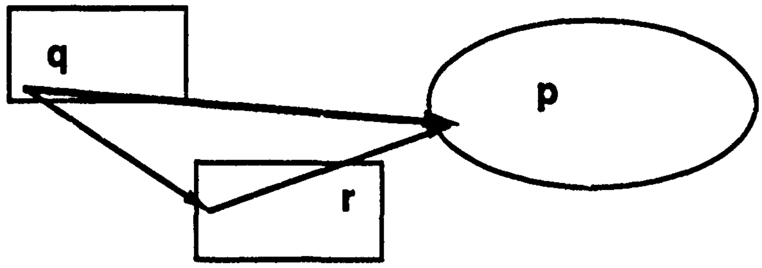

Figure 1: The Generalized Modus Ponens.

concept that plays, with respect to the notion of possibility, the function that is fulfilled by the degree of implication function with respect to necessity. In those derivations, by introduction of sharper bounds for certain conditional concepts, we will also be able to improve the quality of the bounds provided by generalized modus ponens rules while being closer in spirit to its usual fuzzy-logic formulation.

3.3.2 Degree of Consistence

A notion that is dual to that of degree of implication is given by a function that measures the pointwise proximity between pairs of possible worlds from an "optimistic" point of view characterizing the degree by which statements that are true in some worlds *may* apply on others. By contrast, the degree of implication measures the extent by which statements that are true in p-worlds *must* hold in q-worlds.

Definition: The *degree of consistence* of p and q is the value

$$\mathbf{C}(p \mid q) = \sup_{w' \vdash q} \sup_{w \vdash p} S(w, w').$$

An immediate consequence of this definition that $\mathbf{C}(\cdot \mid \cdot)$ is a symmetric function that is increasingly monotonic in both arguments (with respect to the $\Rightarrow$). It is also easy to see that the values of the degree of consistence function are never smaller than the corresponding values of the degree of consistence function,

$$\mathbf{I}(p \mid q) \leq \mathbf{C}(p \mid q),$$

as the amount of stretching required to reach p from some "convenient" q-world is smaller (i.e., higher values of S) than that required to reach p from any q-world. In general, however, the degree of consistence function is not transitive, preventing the statement of a "compatibility" counterpart of the generalized modus ponens rule. Its relationship with the degree of implication function expressed by the expression

$$\mathbf{C}(p \mid q) = \sup_{w' \vdash q} \mathbf{I}(p \mid w') = \sup_{w \vdash p} \mathbf{I}(q \mid w)$$

will permit us, nonetheless, to derive a useful bound-propagation expression.

4 POSSIBILITY AND NECESSITY DISTRIBUTIONS

This section presents interpretations of the major constructs of fuzzy logic —possibility and necessity distributions—in terms of similarity-based structures. Possibility and necessity distributions are functions that measure the proximity of either all or some of the worlds in the evidential set to worlds in other sets that are employed as reference landmarks.

The role played by possibility and necessity distributions is similar to that performed by lower and upper bounds of probability distributions (or by the belief and plausibility functions of the Dempster-Shafer calculus of evidence) with respect to probability distributions. The essential difference between these bounds and those provided by possibility/necessity pairs lies in the fundamentally dissimilar character of what is being bound—metric structures relating pairs of worlds in one case; measures of set size, on the other. Furthermore, in the model of possibilistic structures that is presented in this note necessity (possibility) distributions are any lower (upper) bounds of certain metric functions rather than its "best" or "sharpest" bounds. The operations of fuzzy logic allow computation of bounds for some of these measures as a function of bounds of other measures.

4.1 Inverse of a Triangular Norm

When working in ordinary metric spaces, it is often convenient to express the conventional statement of the triangular inequality, i.e.,

$$\delta(w, w') \leq \delta(w, w'') + \delta(w'', w'),$$

in the equivalent form

$$\delta(w, w') \geq |\, \delta(w, w'') - \delta(w', w'')\,|,$$

which utilizes a form of inverse (i.e., the substraction operator $-$) of the function used to express the original inequality (i.e., the addition operator $+$). This notion of inverse may be directly generalized [37] to provide us with the tools required to define possibility and necessity functions and to derive useful forms of the generalized modus ponens involving either type of these constructs.

Definition: If $\circledast$ is a triangular norm, its *pseudoinverse* $\oslash$ is the function defined over pairs of numbers in the unit interval of the real line, by the expression

$$a \oslash b = \sup\{\, c : \; b \circledast c \leq a \,\}.$$

From this definition it is clear that $a \oslash b$ is nondecreasing in a and nonincreasing in b. Furthermore, $a \oslash 0 = 1$ and $a \oslash 1 = a$ for any a in $[0, 1]$. Other important properties of the pseudoinverse function are given in the works of Schweizer and Sklar [37], Trillas and Valverde [43], and Valverde [44].

Examples of the pseudoinverses of important triangular norms are given in Table 1 together with the corresponding conorms.

Table 1: Triangular Norms, Conorms, and Pseudoinverses

Name	T-Norm $a \circledast b$	Conorm $a \oplus b$	Pseudoinverse $a \oslash b$
Lukasïewicz	$\max(a+b-1, 0)$	$\min(a+b, 1)$	$\min(1+a-b, 1)$
Product	ab	$a+b-ab$	a/b, if $b > a$ 1, otherwise
Zadeh	$\min(a, b)$	$\max(a, b)$	a, if $b > a$ 1, otherwise

4.2 Unconditioned Necessity Distributions

We introduce first a family of functions that bound by below the value of the similarity between any evidential world in $\mathcal{E}$ to some world where another proposition p is true. These unconditioned *necessity* distributions are lower bounds for values of the degree of implication $\mathbf{I}(p \mid \mathcal{E})$, which measures the extent by which statements that are true in a reference set (i.e., the subset of p-worlds) must hold in the evidential set.

As observed before, whenever $\mathbf{I}(p \mid \mathcal{E}) = 1$, it is true, under minimal assumptions, that the evidential subset $\mathcal{E}$ is a subset of the set of all p-worlds, or that p necessarily holds in $\mathcal{E}$. If, on the other hand, $\mathbf{I}(p \mid \mathcal{E}) = \alpha < 1$, then p must be stretched a certain amount—with smaller α corresponding to larger stretching—in order for one of its neighborhoods to encompass $\mathcal{E}$.

Definition: If $\mathcal{E}$ is an evidential set, then a a function $\mathbf{Nec}(\cdot)$ defined over propositions in the language $\mathcal{L}$ is called an *unconditioned necessity distribution* for $\mathcal{E}$ if

$$\mathbf{Nec}(p) \leq \mathbf{I}(p \mid \mathcal{E}).$$

4.3 Unconditioned Possibility Distributions

The dual counterpart of the unconditioned necessity distribution is provided by upper bounds of the degree of consistence $\mathbf{C}(p \mid \mathcal{E})$. Whenever $\mathbf{C}(p \mid \mathcal{E}) = 1$, it is easy to see that, under minimal assumptions, there exists a p-world w that is in the evidential set $\mathcal{E}$ or, equivalently, that p (for all we know) is possibly true. If, on the other hand, $\mathbf{C}(p \mid \mathcal{E}) = \alpha < 1$, then there exists a neighborhood (of "size" α) of some p-world that intersects the evidential set.

Definition: If $\mathcal{E}$ is an evidential set, then a function $\mathbf{Poss}(\cdot)$ defined over propositions in the language $\mathcal{L}$ is called an *unconditioned possibility distribution* for $\mathcal{E}$ if

$$\mathbf{Poss}(p) \geq \mathbf{C}(p \mid \mathcal{E}).$$

Since the value $\mathbf{Poss}(p)$ of any possibility function $\mathbf{Poss}(\cdot)$ is an upper bound of the value $\mathbf{C}(p \mid \mathcal{E})$ of the degree of consistence, while the corresponding value $\mathbf{Nec}(p)$ of any necessity function $\mathbf{Nec}(\cdot)$ is a lower bound of $\mathbf{I}(p \mid q)$, it follows that values of a possibility function can never be smaller than the corresponding values of any necessity function, i.e., that

$$\mathbf{Nec}(p) \leq \mathbf{Poss}(p).$$

4.4 Properties of Possibility and Necessity Distributions

In this subsection we will develop similarity-based interpretations for some basic formulæ of possibilistic calculus. These expressions may be thought of as mechanisms that allow the extension of a partially known possibility distribution. For example, the property that

$$\max(\mathbf{Poss}(p), \mathbf{Poss}(q)) \geq \mathbf{C}(p \vee q \mid \mathscr{E}),$$

which is proved below, is the similarity interpretation of the standard rule that allows computation of the value of the possibility value of a disjunction in fuzzy logic, i.e.,

$$\mathbf{Poss}(p \vee q) = \max(\mathbf{Poss}(p), \mathbf{Poss}(q)).$$

Theorem: If p and q are propositions, and if the quantities $\mathbf{Poss}(p), \mathbf{Poss}(q), \mathbf{Nec}(p)$, and $\mathbf{Nec}(q)$ are such that

$$\begin{aligned} \mathbf{Nec}(p) &\leq \mathbf{I}(p \mid \mathscr{E}), & \mathbf{Nec}(q) &\leq \mathbf{I}(q \mid \mathscr{E}), \\ \mathbf{Poss}(p) &\geq \mathbf{C}(p \mid \mathscr{E}), & \mathbf{Poss}(q) &\geq \mathbf{C}(q \mid \mathscr{E}), \end{aligned}$$

then the following statements *(similarity-based interpretations of the basic laws of fuzzy logic)* are valid:

$$\begin{aligned} \max(\mathbf{Nec}(p), \mathbf{Nec}(q)) &\leq \mathbf{I}(p \vee q \mid \mathscr{E}), \\ \max(\mathbf{Poss}(p), \mathbf{Poss}(q)) &\geq \mathbf{C}(p \vee q \mid \mathscr{E}), \\ \min(\mathbf{Poss}(p), \mathbf{Poss}(q)) &\geq \mathbf{C}(p \wedge q \mid \mathscr{E}). \end{aligned}$$

Proof: Note first that since $\mathbf{C}(\cdot \mid \cdot)$ is nondecreasing (with respect to the $\Rightarrow$ order) in its arguments, it is true that

$$\begin{aligned} \mathbf{Poss}(p) &\geq \mathbf{C}(p \mid \mathscr{E}) \geq \mathbf{C}(p \wedge q \mid \mathscr{E}), \\ \mathbf{Poss}(q) &\geq \mathbf{C}(q \mid \mathscr{E}) \geq \mathbf{C}(p \wedge q \mid \mathscr{E}), \end{aligned}$$

whenever $p \wedge q$ is satisfiable, from which it is easy to see that

$$\min(\mathbf{Poss}(p), \mathbf{Poss}(q)) \geq \mathbf{C}(p \wedge q \mid \mathscr{E}).$$

The corresponding result is obvious when $p \wedge q$ is nonsatisfiable.

A similar argument shows, for necessity functions, that

$$\max(\mathbf{Nec}(p), \mathbf{Nec}(q)) \leq \mathbf{I}(p \vee q \mid \mathscr{E}).$$

To prove the disjunctive law for possibilities, notice that if f is any function mapping elements of a general domain D into real numbers, then

$$\sup\{ f(d) : d \in A \cup B \} = \max\left[\sup\{ f(d) : d \in A \}, \sup\{ f(d) : d \in B \}\right].$$

From this equalitv, it is easy to see that if $\mathbf{Poss}(p)$ and $\mathbf{Poss}(q)$ are upper bounds of $\mathbf{I}(p \mid \mathscr{E})$ and $\mathbf{I}(q \mid \mathscr{E})$, respec .vely, then

$$\max(\, \mathbf{Poss}(p), \mathbf{Poss}(q)\,) \geq \mathbf{C}(p \vee q \mid \mathscr{E}),$$

completing the proof of the theorem. ∎

Note, however, that another law commonly given as an axiom for necessity functions does not hold valid in our interpretation. As illustrated in Figure 2, the distance from a point to the intersection of two sets may be strictly larger than the distances to either set (i.e., the similarity will be strictly smaller). In general, therefore, it is

$$\min(\mathbf{Nec}(p), \mathbf{Nec}(q)) \not\leq \mathbf{I}(p \wedge q \mid \mathscr{E}),$$

making invalid, under this interpretation, the conjunctive law for necessities [11]

$$\mathbf{Nec}(p \wedge q) = \min(\mathbf{Nec}(p), \mathbf{Nec}(q)).$$

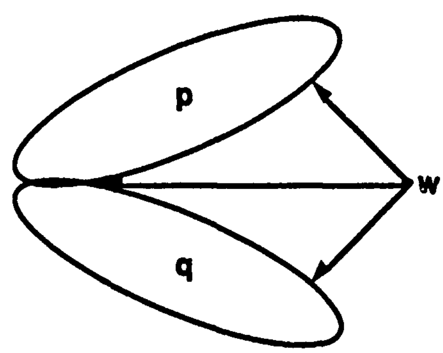

Figure 2: Failure of Conjunctive Necessity.

We may also note in this regard that the similarity-based model that is discussed here does not make use of the notion of negation either as a mechanism to generate dual concepts or on its own right as an important logical concept. It is the intent of the author to study, in the immediate future, alternative models where notions of negation and maximal dissimilarity play more substantive roles.

4.5 Conditional Possibilities and Necessities

The concepts of conditional possibility and necessity are closely related to the previously introduced unconditioned structures. These structures may be thought of as a characterization of the proximity of a world w to some or all of the worlds where a proposition p is true, *given that w is similar in the degree* 1 *to the evidential set* $\mathscr{E}$ (i.e. $w \vdash \mathscr{E}$). With this fact, in mind, we could have used the somewhat baroque formulation

$$\mathbf{C}(p \mid \mathscr{E}) = \sup_{w \vdash \mathscr{E}} \left[\mathbf{I}(p \mid w) \oslash \mathbf{I}(\mathscr{E} \mid w)\right]$$

to define unconditioned possibility distributions—a rather unnecessary effort if we consider that $\mathbf{I}(\mathscr{E} \mid w) = 1$ whenever $w \vdash \mathscr{E}$, showing its obvious equivalence to the simpler form used in Section 3.3.2 above. In spite of such observation, the above identity is important in understanding the purpose of the definitions given below. Those definitions interpret conditional possibilities and necessities as a measure of the proximity of worlds on the evidential set $\mathscr{E}$ to (some or all) worlds satisfying a (conditioned) proposition p relative to their proximity to (some or all) the worlds that satisfy another (conditioning) proposition q.

The mechanism used to specify that relationship, which is closely related in spirit to results of Valverde [44] on the structure of indistinguishability relations, is based on the pseudoinverse function introduced in Section 4.1. The basic idea used by these definitions is also illustrated in Figure 3, where, from the perspective of the evidential world w, the similarity between the p-world u and the q-world v is estimated by means of an inequality that generalizes the "absolute value" form of the triangular inequality, i.e.,

$$\delta(u,v) \geq |\, \delta(u,w) - \delta(v,w) \,|\,,$$

to its similarity-based form

$$S(u,v) \leq \min\left[\, S(u,w) \oslash S(v,w),\; S(v,w) \oslash S(u,w) \,\right].$$

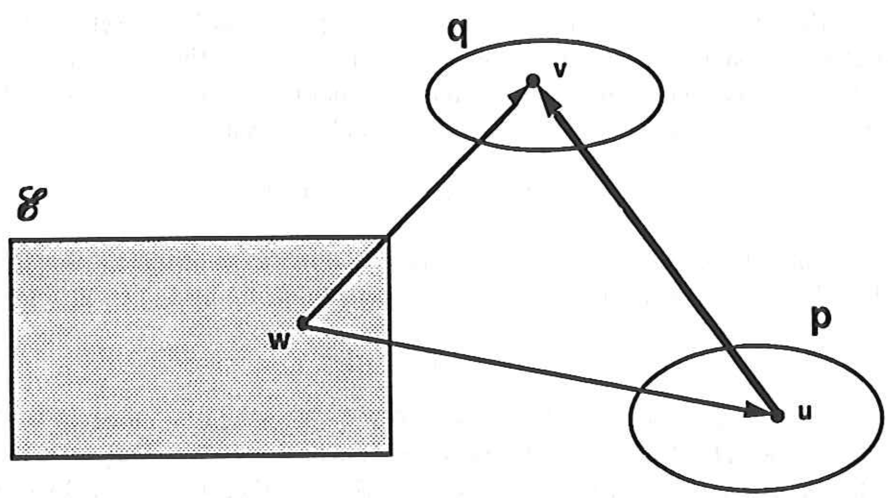

Figure 3: Similarities as Viewed from the Evidential Set.

The required interplay between similarities to conditioning and conditioned sets is captured by the following definitions.

Definition: Let $\mathscr{E}$ be an evidential set. A function $\mathbf{Nec}(\cdot|\cdot)$ mapping pairs of propositions in the language $\mathscr{L}$ into $[0,1]$ is called a *conditional necessity distribution* for $\mathscr{E}$ if

$$\mathbf{Nec}(q|p) \leq \inf_{w \vdash \mathscr{E}} \left[\, \mathbf{I}(q \mid w) \oslash \mathbf{I}(p \mid w) \,\right],$$

for any propositions p and q in $\mathscr{L}$.

Definition: Let $\mathscr{E}$ be an evidential set. A function $\mathbf{Poss}(\cdot|\cdot)$ mapping pairs of propositions in the language $\mathscr{L}$ into $[0,1]$ is called a *conditional possibility distribution* for $\mathscr{E}$ if

$$\mathbf{Poss}(q|p) \geq \sup_{w \vdash \mathscr{E}} \left[\mathbf{I}(q\,|\,w) \oslash \mathbf{I}(p\,|\,w)\right],$$

for any propositions p and q in $\mathscr{L}$.

It is easy to see, from these definitions, that the values of a conditional necessity distribution are never larger than the corresponding values of any conditional possibility distribution, i.e.,

$$\mathbf{Nec}(q|p) \leq \mathbf{Poss}(q|p).$$

Furthermore, since $\mathbf{I}(\cdot\,|\,\cdot)$ is $\circledast$-transitive, then

$$\mathbf{I}(q\,|\,w) \geq \mathbf{I}(q\,|\,p) \circledast \mathbf{I}(p\,|\,w).$$

From this inequality and the definition of pseudoinverse of a triangular norm, it is easy to see that any necessity function satisfies the inequality

$$\mathbf{Nec}(q|p) \geq \mathbf{I}(q\,|\,p),$$

i.e., the bounds for necessity functions provided by the evidential-set perspective are stronger than those that can be obtained by direct use of the degree of implication function.[9]

Note also that if $\mathbf{Nec}(p) = 1$, indicating that $\mathbf{I}(p\,|\,\mathscr{E}) = 1$, and if $\mathbf{Nec}(q|p) = 1$, then the above definition of conditional necessity shows that $\mathbf{I}(q\,|\,\mathscr{E}) = 1$, indicating that $\mathbf{Nec}(q)$ may be taken to be equal to 1, thus generalizing the well-known axiom (consequential closure) of certain modal systems (e.g., the system $\mathbf{T}$, as discussed in Hughes and Creswell [21])

$$\text{If } \mathbf{N}p \text{ and } \mathbf{N}(p \rightarrow q), \text{ then } \mathbf{N}q.$$

The definitions above can also be further interpreted as a way to compare the similarities between evidential worlds and those in the conditioning and conditioned sets by noting that whenever

$$\mathbf{I}(q\,|\,w) \geq \mathbf{I}(p\,|\,w),$$

for every evidential world $w \vdash \mathscr{E}$, then $\mathbf{Nec}(q|p)$ may be chosen to be equal to 1. Similarly, if there exists some world $w \vdash \mathscr{E}$ where this inequality holds, then it is $\mathbf{Poss}(q|p) = 1$. In either case, however, the maximum value for the conditional distribution (i.e., 1) is reached when the proximity of one evidential world w—in the case of possibilities—or of every one of them—in the case of necessities—to a world w_q in the conditioned set exceeds the proximity of w to the conditioning set p. In either case, once again recurring to an apparent notational overkill, we may state this fact by means of the identity function τ in the unit interval:

$$\tau\colon [0,1] \mapsto [0,1] : \alpha \mapsto \alpha,$$

in the form

$$\mathbf{I}(q\,|\,w) \geq \tau\big(\mathbf{I}(p\,|\,w)\big),$$

[9] A dual inequality for possibilities involving $\mathbf{C}(q\,|\,p)$ does not hold in general. It is easy to see, however, that $\mathbf{C}(q\,|\,\mathscr{E}) \oslash \mathbf{I}(p\,|\,\mathscr{E})$ is a possibility function for q given p.

for some $w \vdash \mathscr{S}$ in the case of possibilities, with the same inequality holding for every $w \vdash \mathscr{S}$ in the case of necessities. We may, however, conceive of other functions

$$\gamma\colon [0,1] \mapsto [0,1]\colon \alpha \mapsto \gamma(\alpha)\,,$$

with $\gamma(\alpha) \geq \alpha$ to specify a stronger form of implication, as illustrated in Figure 4, i.e.,

$$\mathrm{I}(q \mid w) \geq \gamma\big(\mathrm{I}(p \mid w)\big)\,.$$

Similarly, one may also conceive of functions ψ with $\psi(\alpha) \leq \alpha$ that may be used to model weaker forms of implication.

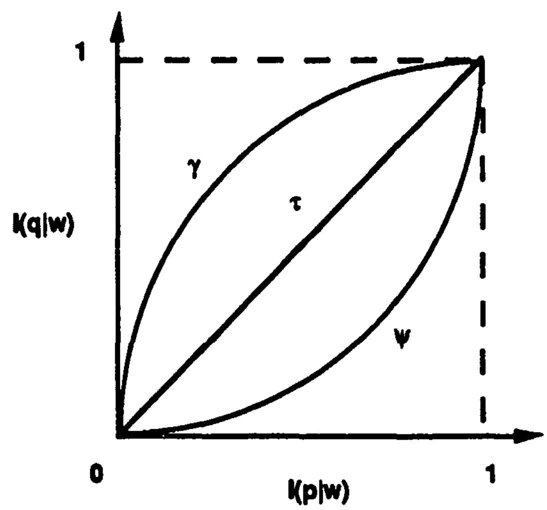

Figure 4: Examples of Possible Similarity Relationships between Conditioning and Conditioned Sets.

Possibilistic calculi based on the propagation of truth-mappings of this type, first proposed by Baldwin [2], are utilized in the RUM [4,5] and MILORD [18] expert systems. The particular case when $\gamma = \tau$, stating that every α-cut of the conditioning proposition p is fully enclosed (in the conventional sense) in the α-cut of the conditioned proposition q, has been called the *truth* mapping in the fuzzy logic literature.

The primary purpose of conditional distributions, however, is to provide a quantitative measure of the strength by which one proposition may be said to imply another with a view to extend inferential procedures by means of structures that superimpose the topological notion of continuity upon a logical framework concerned with propositional validity.

5 GENERALIZED INFERENCE

The major inferential tool of fuzzy logic is the *compositional rule of inference* of Zadeh [53], which generalizes the corresponding classical rule of inference by its ability to infer valid statements even when a perfect match between facts and rule antecedent does not exist, i.e.,

$$\text{from} \quad \frac{\begin{array}{l} p \\ p \;\rightarrow\; q \end{array}}{q} \quad \text{to its "approximate" version} \quad \frac{\begin{array}{l} p' \\ p \;\rightarrow\; q \end{array}}{q'} \;,$$

where p' and q' are similar to p and q, respectively. In this sense, the generalized modus ponens operates as an "interpolation" (or, more precisely, as an "extrapolation") procedure in possible-world space.

Unlike the interpolation procedures of numerical analysis, however, which yield estimates of function value, this extrapolation procedure approximates truth in the sense that it produces a proposition that is both more general than the consequent of the inferential rule and resembles it to some degree (which is a function of the degree by which p' resembles p). The "extrapolated conclusion," however, is a correctly derived proposition, i.e., the result of a sound logical procedure rather than of an approximate heuristic technique.

5.1 Generalized Modus Ponens

The theorems that are proven below are based on the use of a family $\mathscr{P}$ of propositions that partitions the universe of discourse $\mathcal{U}$ in the sense that every possible world will satisfy at least one proposition in $\mathscr{P}$.

Definition: If $\mathscr{P}$ is a subset of satisfiable propositions in $\mathscr{L}$ such that if w is a possible world in the universe $\mathcal{U}$, then there exists a proposition p in $\mathscr{P}$ such that $w \vdash p$, then the family $\mathscr{P}$ is called a *partition* of $\mathcal{U}$.

These results make use of information such as the values of the unconditioned necessity (resp., possibility) distributions for antecedent propositions p in the family $\mathscr{P}$ together with the values $\mathbf{Nec}(q|p)$ (resp., $\mathbf{Poss}(q|p)$) to "extend" the unconditioned distributions to the "consequent" proposition q. In this sense, these findings interpret, in the same spirit used in the theorem of Section 4.4 for other basic laws, the generalized modus ponens laws of fuzzy logic:

$$\begin{aligned} \mathbf{Nec}(q) &= \sup_{\mathscr{P}} \left[\mathbf{Nec}(q|p) \circledast \mathbf{Nec}(p) \right], \\ \mathbf{Poss}(q) &= \sup_{\mathscr{P}} \left[\mathbf{Poss}(q|p) \circledast \mathbf{Poss}(p) \right]. \end{aligned}$$

Theorem *(Generalized Modus Ponens for Necessity Functions)*: Let $\mathscr{P}$ be a partition of $\mathcal{U}$ and let q be a proposition. If $\mathbf{Nec}(p)$ and $\mathbf{Nec}(q|p)$ are real values, defined for every proposition p in the partition $\mathscr{P}$, such that

$$\begin{aligned}\mathbf{Nec}(p) &\leq \mathbf{I}(p \mid \mathscr{E}),\\ \mathbf{Nec}(q|p) &\leq \inf_{w \vdash \mathscr{E}} \left[\mathbf{I}(q \mid w) \oslash \mathbf{I}(p \mid w)\right],\end{aligned}$$

then the following inequality is valid

$$\sup_{\mathscr{P}} \left[\mathbf{Nec}(q|p) \circledast \mathbf{Nec}(p)\right] \leq \mathbf{I}(q \mid \mathscr{E}).$$

Proof: Note first that since $\oslash$ is nonincreasing in its second argument and since

$$\mathbf{I}(p \mid \mathscr{E}) \leq \mathbf{I}(p \mid w)$$

for every evidential world w, it is

$$\mathbf{Nec}(q|p) \leq \inf_{w \vdash \mathscr{E}} \left[\mathbf{I}(q \mid w) \oslash \mathbf{I}(p \mid w)\right] \leq \inf_{w \vdash \mathscr{E}} \left[\mathbf{I}(q \mid w) \oslash \mathbf{I}(p \mid \mathscr{E})\right].$$

It follows then from the monotonicity and continuity of $\circledast$ with respect to its arguments that

$$\begin{aligned}\mathbf{Nec}(p) \circledast \mathbf{Nec}(q|p) &\leq \mathbf{I}(p \mid \mathscr{E}) \circledast \inf_{w \vdash \mathscr{E}} \left[\mathbf{I}(q \mid w) \oslash \mathbf{I}(p \mid \mathscr{E})\right]\\ &= \inf_{w \vdash \mathscr{E}} \left[\mathbf{I}(p \mid \mathscr{E}) \circledast \left(\mathbf{I}(q \mid w) \oslash \mathbf{I}(p \mid \mathscr{E})\right)\right]\\ &\leq \inf_{w \vdash \mathscr{E}} \mathbf{I}(q \mid w)\\ &= \mathbf{I}(q \mid \mathscr{E})\end{aligned}$$

since

$$\mathbf{I}(p \mid \mathscr{E}) \circledast \left(\mathbf{I}(q \mid w) \oslash \mathbf{I}(p \mid \mathscr{E})\right) \leq \mathbf{I}(q \mid w),$$

because of the definition of $\oslash$ and the continuity of $\circledast$.

Since the above inequality is valid for any proposition p in $\mathscr{P}$, the theorem follows. ∎

A dual result also holds for possibility functions.

Theorem *(Generalized Modus Ponens for Possibility Functions)*: Let $\mathscr{P}$ be a partition of $\mathcal{U}$ and let q be a proposition. If $\mathbf{Poss}(p)$ and $\mathbf{Poss}(q|p)$ are real values, defined for every proposition p in $\mathscr{P}$, such that

$$\begin{aligned}\mathbf{Poss}(p) &\geq \mathbf{C}(p \mid \mathscr{E}),\\ \mathbf{Poss}(q|p) &\geq \sup_{w \vdash \mathscr{E}} \left[\mathbf{I}(q \mid w) \oslash \mathbf{I}(p \mid w)\right],\end{aligned}$$

then the following inequality is valid

$$\sup_{\mathscr{P}} \left[\mathbf{Poss}(q|p) \circledast \mathbf{Poss}(p)\right] \geq \mathbf{C}(q \mid \mathscr{E}).$$

Proof: Note first that if w is an evidential world, then

$$\mathbf{C}(p \mid \mathscr{E}) \geq \mathbf{I}(p \mid w).$$

It follows then from the nonincreasing nature of $\oslash$ with respect to its second argument that

$$\begin{aligned}\mathbf{Poss}(q|p) &\geq \sup_{w\vdash\mathscr{E}} \left[\mathbf{I}(q\,|\,w)\oslash\mathbf{I}(p\,|\,w)\right]\\ &\geq \sup_{w\vdash\mathscr{E}} \left[\mathbf{I}(q\,|\,w)\oslash\mathbf{C}(p\,|\,\mathscr{E})\right],\end{aligned}$$

and, therefore, that

$$\mathbf{Poss}(q|p)\circledast\mathbf{Poss}(p) \geq \sup_{w\vdash\mathscr{E}} \left[\mathbf{I}(q\,|\,w)\oslash\mathbf{C}(p\,|\,\mathscr{E})\right]\circledast\mathbf{C}(p\,|\,\mathscr{E}).$$

Taking now, in the above expression, the supremum with respect to all propositions p in $\mathscr{P}$, it is

$$\sup_{\mathscr{P}} \left[\mathbf{Poss}(q|p)\circledast\mathbf{Poss}(p)\right] \geq \sup_{\mathscr{P}} \left[\sup_{w\vdash\mathscr{E}} \left[\mathbf{I}(q\,|\,w)\oslash\mathbf{C}(p\,|\,\mathscr{E})\right]\circledast\mathbf{C}(p\,|\,\mathscr{E})\right]. \tag{1}$$

Note, however, that since $\mathscr{P}$ is a partition, there always exists a proposition $\hat{p}$ in $\mathscr{P}$ such that $\mathbf{C}(\hat{p}\,|\,\mathscr{E})=1$ (i.e., $\hat{p}$ "intersects" $\mathscr{E}$) and, therefore,

$$\begin{aligned}\sup_{\mathscr{P}} \left[\sup_{w\vdash\mathscr{E}} \left[\mathbf{I}(q\,|\,w)\oslash\mathbf{C}(p\,|\,\mathscr{E})\right]\circledast\mathbf{C}(p\,|\,\mathscr{E})\right] &\geq \sup_{w\vdash\mathscr{E}} \left[\mathbf{I}(q\,|\,w)\oslash\mathbf{C}(\hat{p}\,|\,\mathscr{E})\right]\circledast\mathbf{C}(\hat{p}\,|\,\mathscr{E})\\ &= \sup_{w\vdash\mathscr{E}} \mathbf{I}(q\,|\,w)\\ &= \mathbf{C}(q\,|\,\mathscr{E}).\end{aligned} \tag{2}$$

The thesis follows at once by combination of the inequalities (1) and (2). ∎

Finally, notice also that, although the theorems above have been characterized as duals, it is not necessary that $\mathscr{P}$ be a partition for the generalized modus ponens for necessities to hold, while the proof of its possibilistic counterpart relies on such assumption. It should be clear, however, that richer propositional collections $\mathscr{P}$ would lead to better lower bounds for values of the degree of implication $\mathbf{I}(q\,|\,\mathscr{E})$.

5.2 Variables

The $\circledast$-transitivity property of $\mathbf{I}$ is the essential fact expressing the relationships between the degrees of implication of three propositions that were proven in the previous section. The statements of these relations in most works devoted to fuzzy logic are made, however, using special subsets of the universe of discourse that are described through the important notion of *variable*. Introduction of this concept, which is also central to other approximate reasoning methodologies, permits us to make a clearer distinction between similarities defined, in some absolute sense, from the joint viewpoint of several respects and related proximity measures that compare objects (in our case, possible worlds) from the marginal viewpoint of one or more variables.

In what follows, we will assume that only certain propositions, specifying the value of a system variable belonging to a finite set

$$\mathscr{V} = \{X, Y, Z, \dots\},$$

will be used to characterize possible worlds.

The propositions of interest are those formed by logical combination of statements of the type

"The value of the variable V is v,"

where V is in the variable set $\mathscr{V}$, and where v is a specific value in the domain $\mathscr{D}(V)$ of the variable V.

We will also assume that, in any possible world, the value of any variable is a member of the corresponding domain of definition of the variable. In the context of our discussion, we will not need to make special assumptions about the scalar or numeric nature of the state variables, using the notion in the same primitive and general sense in which it is customarily used in the predicate calculus.

We will be specially interested in subsets, called *variable-sets*, of the universe $\mathcal{U}$ consisting of worlds where the value of some variable V is equal to a specified value v. We will denote by $[X = x]$ (similarly $[Y = y]$, etc.) the set of all possible worlds where the proposition "The value of the variable X is x" is true. Clearly, the variable-sets in the collection

$$\{\,[X = x] : x \text{ is in } \mathscr{D}(X)\,\}$$

partition the universe into disjoint subsets. These collections have recently been used to characterize the concept of *rough sets* [30], of importance in many information-system analysis problems, including some that arise in the context of approximate reasoning. A similar notion has also been used also to describe algorithms for the combination of probabilities and of belief functions [39].

To simplify the notation we will write

$$w \vdash x,\ w \vdash y, \ldots$$

as shorthand for $w \vdash [X = x]$, $w \vdash [Y = y]$, ... , respectively.

5.2.1 Possibilistic Structures and Laws

The usual statements of the laws of fuzzy logic are made, as mentioned before, through the use of variables rather than by means of general symbolic expressions. It is customary, for example, to speak of the possibility of the variable X taking the value x, to describe the value that a possibility function for an evidential set $\mathscr{E}$ attains for the proposition $[X = x]$.

In our model, we will say therefore, that a function

$$\mathbf{Poss}(\cdot) \colon \mathscr{D}(X) \mapsto [0, 1]$$

is a possibility function for the evidential set $\mathscr{E}$ and the variable X, whenever

$$\mathbf{Poss}(x) \geq \mathbf{C}\left(\,[X = x] \mid \mathscr{E}\,\right),$$

for all values x in the domain $\mathscr{D}(X)$. Similarly, we will say that $\mathbf{Nec}(\cdot)$ is a necessity function for X whenever

$$\mathbf{Nec}(x) \leq \mathbf{I}\left(\,[X = x] \mid \mathscr{E}\,\right),$$

for all values x in $\mathscr{D}(X)$.

If possibility distributions are point functions defined in this way as point functions in the variable domain $\mathscr{D}(X)$, then it is possible to use the disjunctive laws of fuzzy logic proved in Section 4.4 to extend their definition over the power set of $\mathscr{D}(X)$, i.e.,

$$\begin{aligned} \mathbf{Nec}(A \cup B) &= \max\left[\mathbf{Nec}(A), \mathbf{Nec}(B)\right], \\ \mathbf{Poss}(A \cup B) &= \max\left[\mathbf{Poss}(A), \mathbf{Poss}(B)\right], \end{aligned}$$

where A and B are subsets of the domain $\mathscr{D}(X)$. These equations are usually given as the basic disjunctive laws of possibility distributions.

Note that, using such extensions, both possibility and necessity functions are nondecreasing functions (with respect to the order induced by set inclusion). The value of $\mathbf{Nec}(A)$ measures the extent by which the evidence supports the statement that the variable value necessarily lies in the subset A of its domain of definition, with a dual interpretation being applicable for possibility distributions.

5.2.2 Marginal and Joint Possibilities

The original similarity relation introduced in Section 3.1 may be considered to be a measure of proximity between possible worlds from the joint viewpoint of all system variables. The notion of variable permits, however, the definition of similarities from the restricted viewpoint of some variables or subsets of variables.

These restricted perspectives play a role with respect to the original similarity S that is analogous to that of marginal probability distributions with respect to joint probability distributions. To derive useful expressions that describe similarities between two values x and x' of the same variable X, it should be noted first that the degree of implication $\mathbf{I}(\cdot\,|\,\cdot)$ is transitive. This fact permits the application of a theorem of Valverde [44] to define a function S_X by means of the expression

$$S_X\colon \mathscr{D}(X) \times \mathscr{D}(X) \mapsto [0,1]\colon (x, x') \mapsto \min\left[\mathbf{I}(x\,|\,x'), \mathbf{I}(x'\,|\,x)\right].$$

Defined in this way as a "symmetrization" of the *preorder* induced by the degree of implication $\mathbf{I}(\cdot\,|\,\cdot)$, the marginal similarity S_X has the properties of a similarity function. Furthermore, the "projection" operation entailed by the use of $\mathbf{I}(x\,|\,x')$, based on the projection of every x'-world into the set of x-worlds), may be considered to be the basic mechanism to transform the original similarity function into one that only discern differences in the values of the variable X.

It must be noted, however, that, unless additional assumptions are made about the nature of the original similarity S, the function S_X fails to satisfy the intuitive requirement

$$S(w, w') \le S_X(w, w'),$$

whenever $w \vdash x$ and $w' \vdash x'$ i.e., the similarity between two objects from a restricted viewpoint is always higher than their similarity from more general regards that encompass additional criteria of comparison.

Although considerable research remains to identify alternative definitions of marginal similarities that are not hampered by this problem, a basic result of Valverde [44], presented in Section 6.2 below, appears to provide the essential tool that must be employed in to produce the required coarser measures. The role of additional reasonable assumptions that might be demanded from S so as to facilitate the construction of marginal similarities with desirable characteristics is also the object of current investigations of the author.

5.2.3 Conditional Distributions and Generalized Inference

The basic conditional structures of fuzzy logic are usually defined as elastic constraints that restrict the values of a variable given those of another. By simple extension of our previous convention to conditional structures, we will write $\mathbf{Nec}(y|x)$ and $\mathbf{Poss}(y|x)$, as shorthand for

$$\mathbf{Nec}([Y = y] \mid [X = x]) \quad \text{and} \quad \mathbf{Poss}([Y = y] \mid [X = x]),$$

respectively.

If a classical (i.e., Boolean) inferential rule of the type

"If $X = x$, then Y is in $R(x)$"

is thought of as the definition of a relation R defined over pairs (x, y) in the Cartesian product $X \times Y$, then such a relation may be used to define a multivalued mapping that maps possible values of X into possible values of Y as illustrated in Figure 5.

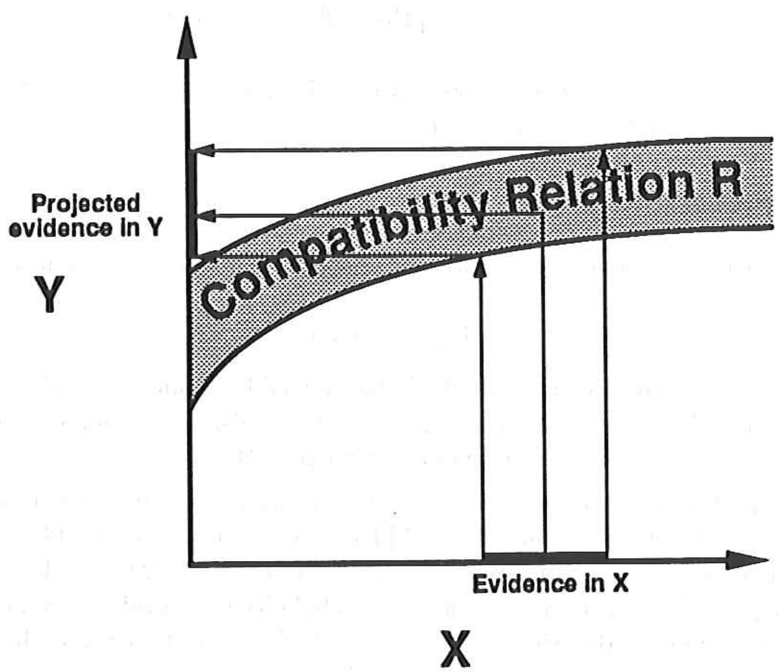

Figure 5: Inference as a Compatibility Relation.

Such a *compatibility relation* perspective was an essential element of the original formulations of both the Dempster-Shafer calculus of evidence [8] where distributions in some space (i.e., the domain of some variable X) are mapped into distributions of another variable (i.e., the domain of another variable Y) by direct transfer of "mass" from individual values to the union of their mapped projections and the compositional rule of inference [51].

Note that, whenever $\mathbf{Poss}(y|x) = 1$, if the bound is actually attained, i.e, if

$$\sup_{w \vdash y} \left[\mathbf{I}(y \mid w) \oslash \mathbf{I}(x \mid w) \right] = 1 ,$$

then it is possible for an evidential world w in $[X = x]$ (i.e., $\mathbf{I}(x \mid w) = 1$) to be such that $w \vdash y$. Pairs (x, y) such that $\mathbf{Poss}(y|x) = 1$ may be considered to approximate the *core*[10] of a generalized inferential relation that allows to determine bounds for the similarity between evidential worlds and those in the variable set $[Y = y]$ on the basis of knowledge of similar bounds applicable to the variable set $[X = x]$. This relation, which is the fuzzy extension of the classical compatibility mapping R illustrated in Figure 5, may be thought as a descriptor of the behavior, for x-worlds, of the values of the variable Y "near" R. The compatibility relation is itself approximated by (or embedded in) the core of the conditional possibility distribution, i.e., worlds w such that $w \vdash x$ and $w \vdash y$, with $\mathbf{Poss}(y|x) = 1$.

Since the collection of the sets $[X = x]$ partitions the universe $\mathcal{U}$ into disjoint sets, then the generalized modus ponens laws may be readily stated in terms of variable values as

$$\begin{aligned} \mathbf{Nec}(y) &= \sup_x \left[\mathbf{Nec}(y|x) \circledast \mathbf{Nec}(x) \right] , \\ \mathbf{Poss}(y) &= \sup_x \left[\mathbf{Poss}(y|x) \circledast \mathbf{Poss}(x) \right] , \end{aligned}$$

clearly showing the basic nature of the inferential mapping as the composition of relational combination (i.e., $\circledast$-"intersection") and projection (i.e., maximization).

5.2.4 Fuzzy Implication Rules

In this section we will examine proposed interpretations for conditional rules, usually stated in the form

$$\text{If } X \text{ is } A, \text{ then } Y \text{ is } B ,$$

within the context of possibilistic logic. While, in two-valued logic, any such rule simply states that whenever a condition A is true, another condition B also holds, various interpretations have been proposed for rules expressing other notions of conditional truth.

In the case of probabilities, for example, degrees of conditionality have been modeled either by means of conditional probability values $\mathbf{Prob}(A \mid B)$, which measure the likelihood of B given the assumed truth of A, or by the alternative interpretation $\mathbf{Prob}(\neg A \vee B)$, used by Nilsson [29] in his probabilistic logic, which esssentially quantifies the probability that a rule is a valid component of a knowledge base. Either one of these interpretations is valid in particular contexts being, respectively, the probabilistic extensions of the so called "de re," i.e.,

$$p \rightarrow \Pi\, q ,$$

and "de dicto", i.e.,

$$\Pi\, (p \rightarrow q) ,$$

interpretations of conditionals in modal logic.

[10] The core of a fuzzy set μ: $\mathcal{U} \mapsto [0, 1]$ is the set of all points w such that $\mu(w) = 1$, i.e., the points that "fully" belong to μ.

In fuzzy logic, two major interpretations have been advanced to translate conditional rules,[11] with A and B corresponding to the fuzzy sets

$$\mu_A\colon X \mapsto [0,1]\,, \quad \text{and} \quad \mu_B\colon Y \mapsto [0,1]\,.$$

The first interpretation was originally proposed by Zadeh [52], as a formal translation of the statement

If μ_A is a possibility for X, then μ_B is a possibility distribution for Y.

This conditional statement, which may be regarded as a constraint on the values of one variable given those of another, states the existence of a conditional possibility function $\mathbf{Poss}(\cdot|\cdot)$ such that

$$\mu_B(y) \geq \sup_x \left[\mathbf{Poss}(y|x) \circledast \mu_A(x)\right] \geq \mathbf{Poss}(y|x) \circledast \mu_A(x)\,.$$

Recalling now the definition and properties of the pseudoinverse, we may restate this particular interpretation as

$$\mathbf{Poss}(y|x) = \mu_B(y) \oslash \mu_A(x) \geq \mathbf{I}(y \mid w) \oslash \mathbf{I}(x \mid w)\,,$$

for every world $w \vdash \mathscr{E}$.

In Zadeh's original formulation, made within the context of a calculus based on the minimum function as the T-norm, conditionals were, however, formally translated by means of the pseudoinverse of the Lukasïewicz T-norm. Certain formal problems associated with such a combination were pointed out by Trillas and Valverde [42], who developed translations consistent with the T-norm used as the basis for the possibilistic calculus.

Using the characterization of conditionals introduced in Section 4.5, this relation may also be thought of as a measure of the degree by which a possibility for Y exceeds a fraction (measured by the conditional possibility distribution) of a given possibility distribution for X. In particular, whenever $\mathbf{Poss}(y|x) = 1$, then $\mu_B(y) \geq \mu_A(x)$, indicating the *possible* existence —since $\mathbf{Poss}(y|x)$ is only an upper bound of $\mathbf{I}(y \mid w) \oslash \mathbf{I}(x \mid w)$ — of an evidential world such that $w \vdash x$ and $w \vdash y$, with x in A and y in B.

As illustrated in Figure 6, where it has been assumed that the underlying metric (i.e., dissimilarity) is proportional to the euclidean distance in the plane, the core of the corresponding conditional possibility distribution is an (upper) approximant of a classical compatibility relation (indicated by the shaded area in the figure) that fans outward from the Cartesian product of the cores of A and B. If this interpretation is taken, whenever several such rules are available, then each one of these rules will lead to a separate possibility distribution. Combination of these upper bounds by minimization results in a sharper possibility estimate that represents the "integrated" effect of the rule set.

The second interpretation of conditional relations, leading to a wide variety of practical applications [41], was utilized by Mamdani and Assilian to develop fuzzy controllers. The basic idea underlying this explanation follows an approach originally outlined by Zadeh [47,48,51]. In this case, a number of conditional statements of the form

$$\text{If } X \text{ is } A_k, \text{ then } Y \text{ is } B_k\,, \qquad k = 1, 2, \ldots, n\,,$$

are given as a combined "disjunctive" description of the relation between X and Y, rather than as a set of independently valid rules. The purpose of this rule set is the approximation of the

[11] A rather encompassing account of potential fuzzy reasoning mechanisms can be found in a paper by Mizumoto, Fukami, and Tanaka. [27]

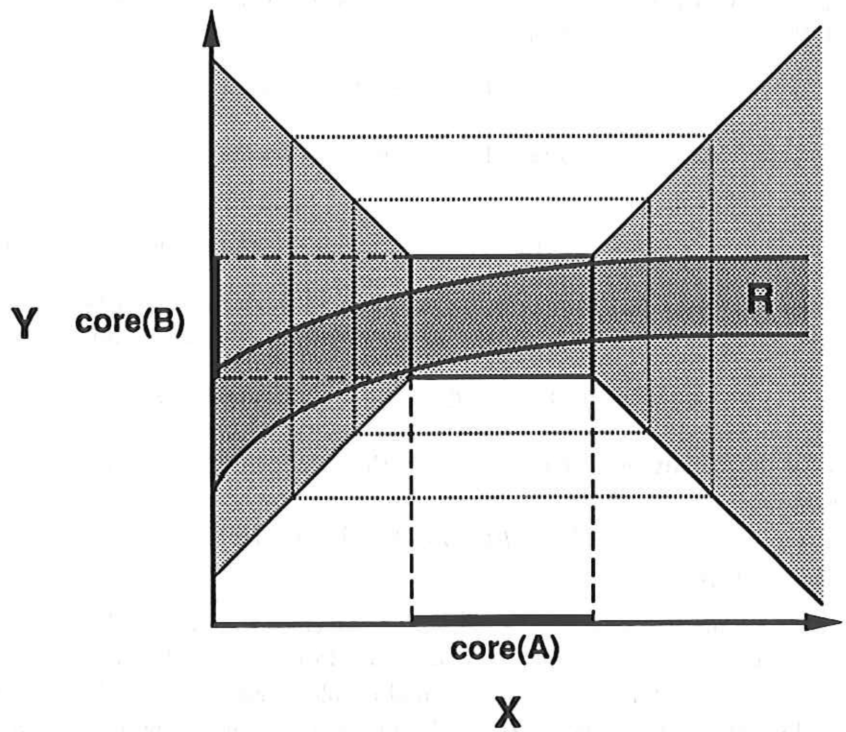

Figure 6: Rules as Possibilistic Approximants of a Compatibility Relation.

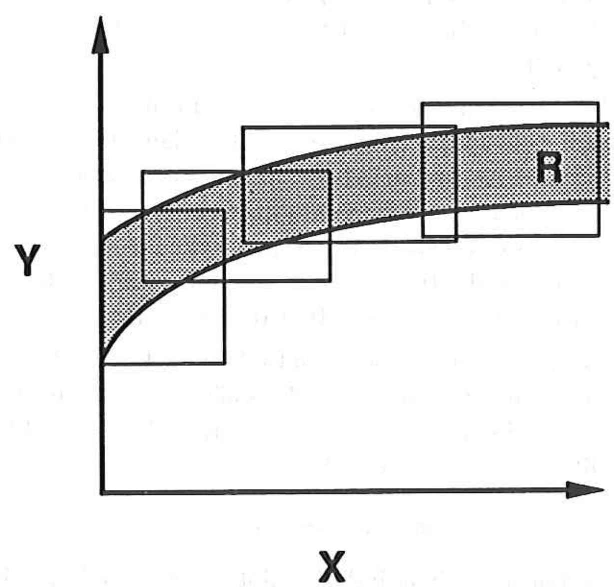

Figure 7: Rule-Sets as Possibilistic Approximants of a Compatibility Relation

compatibility relation by a "fuzzy curve" generated by disjunction of all the rules in the set, as shown in Figure 7.

Recalling the characterization of conditioning as an extension of a classical compatibility relation, we may say that the core of the compatibility relation is approximated by above by the union

$$\bigcup_{k=1}^{n} \left[\operatorname{core}(\mu_{A_k}) \times \operatorname{core}(\mu_{B_k}) \right]$$

of the Cartesian products of the cores of the fuzzy sets for A_k and B_k. In this case the multiple rules are meant to approximate some region of possible (X, Y) values, and the result of application of individual component rules must be combined using maximization to produce a conditional possibility function. We may say, therefore, that under the Zadeh-Mamdani-Assilian (ZMA) interpretation, the function

$$\mathbf{Poss}(y|x) = \sup_k \left[\min(\mu_A(x), \mu_B(y)) \right],$$

is a conditional possibility for Y given X.

It is important to note that the two interpretations of fuzzy rules that we have just examined are based on different approaches to the approximation (by above) of the value

$$\sup_{w \vdash \mathscr{E}} \left[\mathbf{I}(y \mid w) \oslash \mathbf{I}(x \mid w) \right],$$

being, in the the case of the Zadeh-Trillas-Valverde (ZTV) method, the result of the *conjunction* of multiple fuzzy relations such as that illustrated in Figure 8, while, in the case of the ZMA logic, the construction requires *disjunction* of relations such as that illustrated in Figure 9.

The difference between both approaches when combining several rules is illustrated also in Figures 10 and 11, showing the contour plots for the α-cuts of the fuzzy relations that are obtained in a simple example involving four rules. In these figures, the rectangles with a dark outline correspond to the Cartesian products of the cores of the antecedents A_k and B_k. Darker shades of gray correspond to higher degrees of membership.

The reader should be cautioned, however, about the potential for invalid comparisons that may result from hasty examination of these figures. Each formalism should be regarded as a procedure for the approximation of a compatibility relation that is based on a different approach for the description of relationships between variables. In the case of the ZMA interpretation, the intent is to generalize the interpolation procedures that are normally employed in functional approximation. As such, this approach may be said to be inspired by the methodology of classical system analysis. The ZTV approach, by contrast, is a generalization of classical logical formulations and may be regarded, from a relational viewpoint, as a procedure to describe a function as the locus of points that satisfies a set of constraints rather than as a subset of "fuzzy points" of a Cartesian product.

Figures 10 and 11, while showing that the same rule sets would lead to radically different results, should not be considered, therefore, to discredit interpolative approaches as such techniques, proceeding from a different perspective, should normally be based on rule sets that are different from those utilized when rules are thought of as independent constraints.

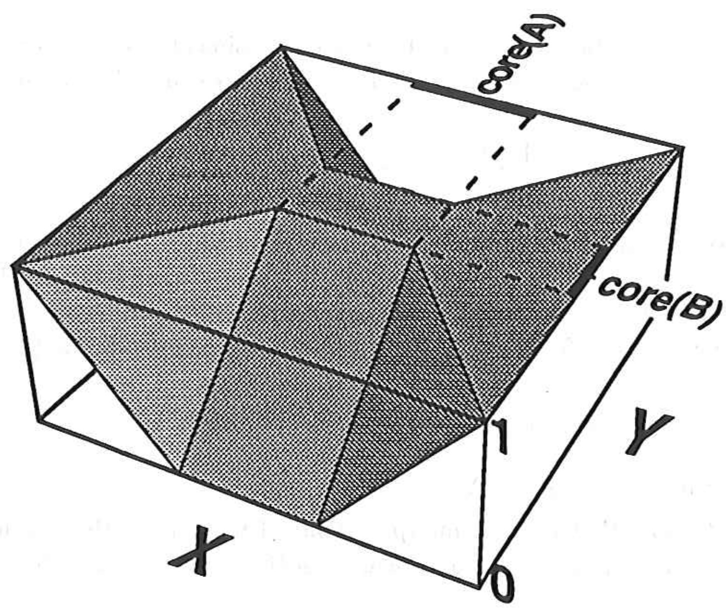

Figure 8: A Possibilistic Conditional Rule (ZTV)

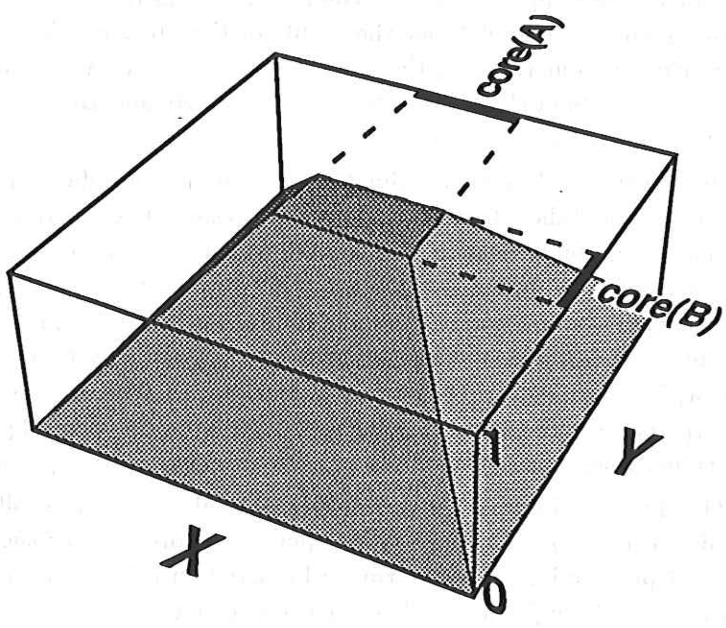

Figure 9: A Component of a Disjunctive Rule Set (ZMA)

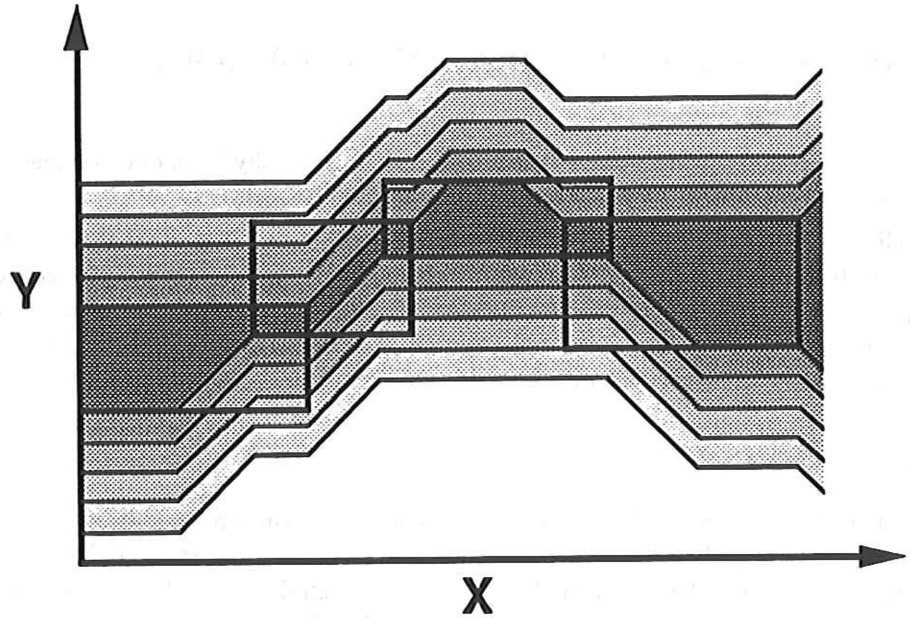

Figure 10: Contour Plots for a Rule Set (ZTV)

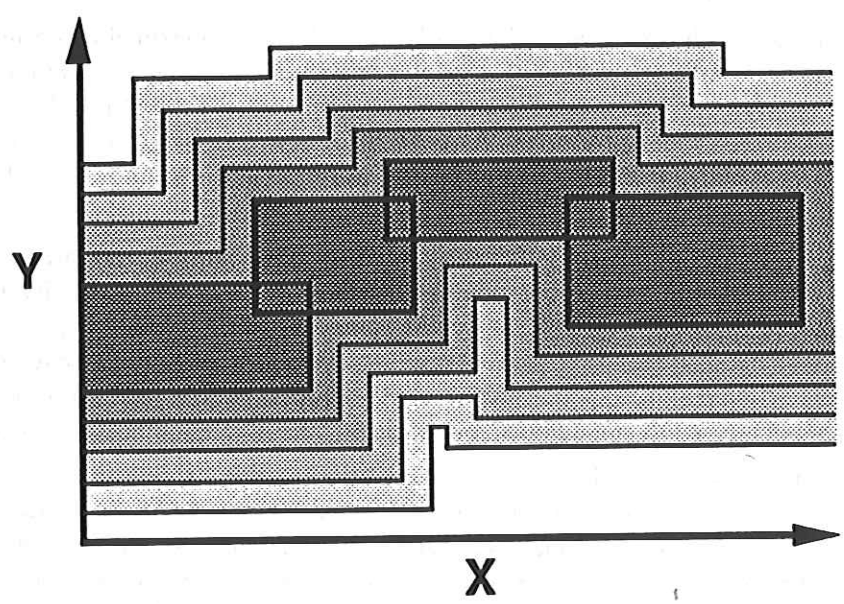

Figure 11: Contour Plots for a Rule Set (ZMA)

6 THE NATURE OF SIMILARITY RELATIONS

In this closing section, we will examine issues that arise naturally from our previous examination of the role of similarities as the semantic bases for possibility theory.

Our discussion focuses on two topics. We look first at the requirements that our theory imposes upon the nature of the scales used to measure proximity or resemblance between possible worlds. Finally, our examination of the interplay between similarities and possibilities turns to issues related to the generation of similarity relations from such sources as domain knowledge that describes significant relations between system variables.

6.1 On Similarity Scales

Our previous interpretation of possibilistic concepts and structures has been based on the use of measures of proximity that quantify interobject resemblance using real numbers between 0 and 1. Our assumptions about the use of the $[0,1]$ interval as a similarity scale have been made primarily, however, as a matter of convenience so as to simplify the description of our model while being consistent with the customary definitions of possibility and necessity distributions as functions taking values in that interval.

Close examination of the actual requirements imposed upon our similarity scales reveals, however, that our measurement domain may be quite general so as to include symbolic structures such as

$$\{\, \textit{identical}, \textit{very similar}, \ldots, \textit{completely dissimilar} \} .$$

Our model is based on the use of a partially ordered set having a maximal and a minimal element that measure identity and complete dissimilarity, respectively. Furthermore, we have assumed the existence of a binary operation (the triangular norm $\circledast$) mapping pairs of possible worlds into real numbers, with certain desirable order-preserving and transitive properties. The concept of triangular norm, however, does not rely substantially on the use of real numbers as its range and may be readily extended to more general partially ordered sets with maximal and minimal elements.

We have also assumed a continuity property for the triangular norm operation. This property, however, simply requires that a notion of proximity also exist among similarity values so as to provide a form of (order-consistent) topology in that space. While, in general, more precise scales will result in more detailed representations of interworld similarity, it is important to stress that the similarity-based model presented here does not rely in "denseness" assumptions such as the existence an intermediate value c between any different values a and b in the similarity-measurement scale.

From a practical viewpoint, the major requirement is to quantify proximity in such a way as to be able to determine that two quantities are similar to some degree (i.e., approximate matching). The degree of precision that such a matching entails is problem-dependent and will be typically the result of conflicting impositions between the desire, on one hand, to keep granularity relatively high to reduce complexity, and the need, on the other, to describe system behavior at an acceptable level of accuracy. The work of Bonissone and Decker [4] is a significant example of the type of systematic study that must be carried out to define similarity scales that are both useful and tractable.

6.2 The Origin of Similarity Functions

The model of fuzzy logic presented in this note is centered on the metric notion of similarity as a primitive concept that is useful to explain the nature of possibilistic constructs and the meaning of possibilistic reasoning. In this formulation, similarities are defined as real functions defined over pairs of possible worlds.

. From this perspective, similarities describe relations of resemblance between objects of high complexity, which, typically, result from consideration of a large number of system variables. Reliance on such complex structures has been the direct consequence of a research program that stressed conceptual clarification as its primary objective. In practice, however, it will be generally difficult to define complex measures that quantify similarity between complex objects on the basis of a large number of criteria.

Similarities provide the framework that is required to understand approximate relations of corelevance, usually stated as generalized conditional rules. The practical generation of similarity functions typically proceeds, however, in the opposite direction, from separate statements about limited aspects of system behavior to general metric structures. Once such resemblance measures are defined, they may be used to express and acquire new laws of system behavior determined, for example, from historical experience with similar systems. Furthermore, such similarity notions may be used as the basis for analogical reasoning systems that try to determine system state on the basis of similarity to known cases [23].

Perhaps the simplest mechanism that may be devised to generate complex metrics from simpler ones is that which starts with measures of resemblance that quantify proximity from a limited viewpoint. These metrics are usually derived, using a variety techniques, in unsupervised pattern classification (or clustering) problems [20]. In many important applications, hierarchical taxonomies —a feature of many representation approaches in artificial intelligence—may be used, often in connection with a variety of weighing schemes—quantifying branching importance—to generate metrics that often satisfy the more stringent requirements of an ultrametric [22].

Classification hierarchies such as those may be thought of as sets of general rules, having a particularly useful structure, that specify interset proximity from relevant, but restricted viewpoints, eventually providing measures of similarity between variable values (i.e., the "leaves" of the taxonomical tree). More generally, however, we may expect that sets of possibilistic rules (i.e., a general knowledge base) defining a general semantic network of corelevance relations may be available as the source for the determination of interobject proximity. These possibilistic semantic networks resemble conventional semantic networks in most regards, being more general in that, in addition to specifying knowledge about system behavior in some subsets of state-space,[12] they also specify characteristics of behavior in neighborhoods of those subsets.

We may think, therefore, that the antecedents of implicational rules define general regions in state space where existence of relevant knowledge may increase insight through application of inferential rules. Using Zadeh's terminology, these antecedents define "granules" that identify important regions of state-space and indicate the level of accuracy that is required (or *granularity*) to perform effective system analysis. In this case, the possibilistic granules correspond to fuzzy sets that are used to specify both what is true in the core of the granule and, with decreasing specificity, what is true in a nested set (i.e., the α-cuts) of its neighborhoods. The ability to specify behavior using such a topological structure results in inferential gains that are the direct consequence of our ability

[12]The expression "state-space" is loosely used here to indicate the space defined by all system variables.

to reason by similarity; an ability that is made possible by the approximate matching property of the generalized modus ponens. From another perspective yet, the fuzzy granules identified by possibilistic rules may also be thought of as generalizations of the arbitrary variable sets used in a variety of artificial intelligence efforts aimed at understanding system behavior using qualitative descriptions of reality [16].

A number of heuristics may be easily formulated to integrate "marginal" measures of resemblance into joint similarity relations. More generally, however, we may state the problem of similarity construction as that of defining metric structures on the basis of knowledge of the aspects of system behavior that are important to its understanding—i.e., the previously mentioned granules, which define what must be distinguished. Since generally those granules are fuzzy sets, the relevance to similarity construction of the following representation theorem, due to Valverde, may be immediately seen:

Theorem [Valverde]: A binary function S mapping pairs of objects of a universe of discourse $\mathcal{U}$ into $[0,1]$ is a similarity relation, if and only if there exists a family $\mathcal{H}$ of fuzzy subsets of $\mathcal{U}$ such that

$$S(w,w') = \inf_{\mathcal{H}} \left[\min \left(h(w) \oslash h(w'),\, h(w') \oslash h(w) \right) \right],$$

for all w and w' in $\mathcal{U}$, where the infimum is taken over all fuzzy subsets h in the family $\mathcal{H}$.

Besides its obvious relevance to the generation of similarity relations from knowledge of important sets in the domain of discourse, Valverde's theorem—resulting originally from studies in pattern recognition—is also of potential significance to the solution of knowledge acquisition problems because of the important relations that exist between learning procedures and structure-discovery techniques such as cluster analysis.

7 CONCLUSION

This note has presented a similarity-based model that provides a clear interpretation of the major structures and methods of possibilistic logic using metric concepts that are formally different from the set-measure constructs of probability theory. Regardless of the potential existence, so far unestablished, of probability-based interpretations for possibilistic structures, this metric model makes clear that there are no compelling reasons to confuse two rather different aspects of uncertainty into a single notion simply because one's favorite theoretical framework, in spite of its otherwise many remarkable virtues, fails to fully capture reality.

Succintly stated, being in a situation that resembles a state of affairs S does not make S likely or viceversa. Furthermore, our reference state may not even be possible in the current circumstances —making it completely unlikely—but we may still find it useful as a comparison landmark.This use of "impossible" examples as a way to illustrate system behavior is very prevalent in human culture, being exemplified by such utterances as "he had the strength of a horse and the swiftness of a swallow," even if it is obvious to all that no such beasts exist other than for such metaphorical purposes.

The insight provided by this model makes it rather obvious that very little can be gained by continuing to assert a potential—although never revealed—encompassing probabilistic interpretation for possibilistic structures that, presumably, would render them unnecessary as serious objects of scientific discourse. In addition, and quite beyond whatever understanding theory may provide, the current success of possibilistic logic as the basis for major systems of important human value [41] —often unmatched by other approaches—should be enough to convince those having more pragmatic perspectives as to its utility.

The task for approximate reasoning researchers is to proceed now beyond unnecessary controversy into the study of the issues that arise from models such as the one presented in this note. Among such questions, further studies of the relations between the notions of possibility, similarity, and negation and of those between probability and possibility are of major importance.

Acknowledgments

The model presented in this note is the product of a long and laborious effort aimed at the explication of possibilistic structures by means of more primitive similarity concepts. During this endeavor, which produced sufficient versions to require eventually a formal numbering system, the author benefitted from the advice and comments of Claudi Alsina, Hamid Berenji, Piero Bonissone, Didier Dubois, Francesc Esteva, Oscar Firschein, Marty Fischler, Pascal Fua, María Angeles Gil, Luis Godo, Andy Hanson, Jerry Hobbs, David Israel, Joan Jacas, Yvan Leclerc, Ramón López de Mántaras, John Lowrance, Abe Mamdani, Bob Moore, Ray Perrault, Henri Prade, Elie Sanchez, Philippe Smets, Tom Strat, Enric Trillas, Llorenç Valverde, Len Wesley, and Lotfi Zadeh. To all of them many thanks.

The basic ideas leading to the model reported in this note were first conceived as part of a basic research effort supported by the United States Air Force Office of Scientific Research. Full development of the similarity-based model was the result of further investigations supported, in addition, by the United States Army Research Office. The author is specially grateful to Dr. Abraham Waksman of the Air Force Office of Scientific Research and Dr. David Hislop of the Army Research Office for their encouragement and for their support of basic research on fundamental issues of approximate reasoning.

The Fulbright Commission for the International Exchange of Scholars, the University of the Balearic Islands, the Center of Advanced Studies of Blanes, the National Research Council of Spain, and the Caixa de Pensions of Barcelona supported several visits by the author to Spanish research centers where he had invaluable exchanges on the subject matter of this technical note.

The assistance of Joani Ichiki, Valerie Maslak, and Diego Ruspini in the preparation of the final manuscript is gratefully acknowledged.

References

[1] C. Alsina and E. Trillas. Additive homogeneity of logical connectives for membership functions. In J.C. Bezdek, editor, *Analysis of Fuzzy Information - Volume I: Mathematics and Logic*, Boca Raton, Florida: CRC Press, 179–184, 1987.

[2] J.F. Baldwin. A new approach to approximate reasoning using a fuzzy logic. *Fuzzy Sets and Systems*, 2:302–335, 1979.

[3] J.C. Bezdek and J.O. Harris. Fuzzy partitions and relations: An axiomatic basis for clustering. *Fuzzy Sets and Systems*, 1:112–127, 1978.

[4] P. Bonissone and K. Decker. Selecting uncertainty calculi and granularity: an experiment in trading-off precision and uncertainty. In L.N. Kanaal and J.F. Lemmer, editors, *Uncertainty in Artificial Intelligence*, Amsterdam: North Holland, 1986.

[5] P.P. Bonissone, S.S. Gans, and K.S. Decker. RUM: a layered architecture for reasoning with uncertainty. In John McDermott, editor, *Proc. Tenth Intern. Joint Conf. on Artificial Intelligence*, 891–896, Los Altos, California: Morgan Kaufmann Publishers, 1987

[6] R. Carnap. *The Logical Foundations of Probability.* Chicago: University of Chicago Press, 1950.

[7] G. Choquet. Théorie des capacités. *Ann. Inst. Fourier (Grenoble)*, V:131–295, 1953.

[8] A.P. Dempster. Upper and lower probabilities induced by a multivalued mapping. *Annals of Statistics*, 38:325–339, 1967.

[9] J. Dieudonné. *Foundations of Modern Analysis.* New York: Academic Press, 1960.

[10] D. Dubois and H. Prade. Fuzzy logics and the generalized modus ponens revisited. *Int. J. of Cybernetic and Systems*, 15:293–331, 1984.

[11] D. Dubois and H. Prade. *Possibility Theory: an Approach to the Computerized Processing of Uuncertainty.* New York: Plenum Press, 1988.

[12] D. Dubois and H. Prade. In search of a modal system for possibility theory. In *Proc. 8th. European Conf. on Artificial Intelligence*, 501–506, Munich: Technical University, 1988.

[13] D. Dubois and H. Prade. An introduction to possibility theory and fuzzy logics. In P. Smets, E.H. Mamdani, D. Dubois, H. Prade, editors, *Non-standard Logics for Automated Reasoning*, 287–326, New York: Academic Press, 1988.

[14] D. Dubois and H. Prade. Representation and combination of uncertainty with belief and possibility measures. *Computational Intelligence,* to appear.

[15] D. Dubois and H. Prade. Rough fuzzy sets and fuzzy rough sets,. *Int. J. General Systems*, to appear.

[16] K. Forbus. Qualitative process theory. *Artificial Intelligence*, 24: 85–168, 1984.

[17] B. Gaines. Fuzzy and probability uncertainty logics. *Inf. Control*, 38:154–169, 1978.

[18] L.I. Godo, R. López de Mántaras, C. Sierra, A. Verdaguer. Managing linguistically expressed uncertainty in MILORD: Application to Medical Diagnosis. Research Report No. 87/2, Group of Logic and Artificial Intelligence, Center of Advanced Studies of Blanes, Blanes, Spain, 1987.

[19] P.R. Halmos. *Measure Theory.* New York: Springer-Verlag, 1974.

[20] J. Hartigan. *Clustering Algorithms*. New York: John Wiley and Sons, 1975.

[21] G.E. Hughes and M.J. Creswell. *An Introduction to Modal Logic*. New York: Methuen, 1972.

[22] N. Jardine and R. Sibson. *Mathematical Taxonomy*. New York: John Wiley and Sons, 1971.

[23] J. Kolodner, editor. Case-Based Reasoning. *Proceedings of a Workshop, Clearwater Beach, Florida, May 1988.* San Mateo, California: Morgan Kaufmann Publishers, 1988.

[24] G. Lakoff. Hedges: A study in meaning criteria and the logic of fuzzy concepts. *J. Philos. Logic*, 2: 458–508, 1973.

[25] D. Lewis. *Counterfactuals*. Cambridge, Massachusetts: Harvard University Press, 1973.

[26] E.H. Mamdani and S. Assilian. An experiment in linguistic synthesis with a fuzzy logic controller. *Int. J. Man-Machine Studies*, 7:1–13, 1975.

[27] M. Mizumoto, S. Fukami, K. Tanaka. Some methods of fuzzy reasoning. In M.M. Gupta, R.K. Ragade, R.R. Yager, editors, *Advances in Fuzzy Set Theory and Applications*, Amsterdam: North Holland, 117–136, 1979.

[28] T. Murai, M. Miyakoshi, M. Shimbo. Fuzzifications of modal operators from the standpoint of fuzzy semantics. In *Proc. 2nd. Int. Fuzzy Systems Assoc. Congress*, 430–432, Tokyo, 1987.

[29] N.J. Nilsson. Probabilistic logic. *Artificial Intelligence*, 28:71–87, 1987.

[30] Z. Pavlak. Rough sets. *Int. J. Comput. Inf. Sci.*, 11:341–356, 1982.

[31] N. Rescher. *Many Valued Logic*. New York: McGraw-Hill, 1969.

[32] E.H. Ruspini. A new approach to clustering. *Inf. Control*, 15:22–32, 1969.

[33] E.H. Ruspini. *A Theory of Cluster Analysis*. Ph.D. Thesis, Department of System Science, University of California, Los Angeles, 1977.

[34] E.H. Ruspini. Recent developments in fuzzy clustering. In R.R. Yager, editor, *Fuzzy Set and Possibility Theory: Recent Developments*, 133–147, New York: Pergamon Press, 1982.

[35] E.H. Ruspini. *The Logical Foundations of Evidential Reasoning*. Technical Note No. 408, Artificial Intelligence Center, SRI International, Menlo Park, California, 1987.

[36] P.K. Schocht. Fuzzy modal logic. In *Proc. Fifth Intern. Sympo. Multiple-Valued Logic*, 176–182, IEEE, 1975.

[37] B. Schweizer and A. Sklar. Associative functions and abstract semigroups. *Publ. Math. Debrecen*, 10:69–81, 1963.

[38] G. Shafer. *A Mathematical Theory of Evidence*. Princeton, New Jersey: Princeton University Press, 1976.

[39] G. Shafer, P.P. Shenoy and K. Mellouli. Propagating belief functions in qualitative Markov trees. *Int.J. Approximate Reasoning*, 1:349–400, 1987.

[40] R.R. Sokal and P.H.A. Sneath. *Principles of Numerical Taxonomy*. San Francisco: Freeman, 1963.

[41] M. Sugeno. *Industrial Applications of Fuzzy Control*. Amsterdam: North Holland, 1985.

[42] E. Trillas and L. Valverde. On some functionally expressible implications for fuzzy set theory. In E.P. Klement, editor, *Proceedings Third International Seminar on Fuzzy Set Theory*, Linz, Austria: Johaness Kepler Univ., 173–190, 1981.

[43] E. Trillas and L. Valverde. On mode and implication in approximate reasoning. In M.M. Gupta, A. Kandel, W. Bandler, J.B. Kiszka, editors, *Approximate Reasoning and Expert Systems*, Amsterdam: North Holland, 157–166, 1985.

[44] L. Valverde. On the structure of F-indistinguishability operators. *Fuzzy Sets and Systems*, 17:313–328, 1985.

[45] A. Tversky. Features of similarity. *Psychological Review*, 84:433–460, 1977.

[46] L.A. Zadeh. Fuzzy sets. *Information and Control*, 8: 338–353, 1965.

[47] L.A. Zadeh. A rationale for fuzzy control. *Journal of Dynamic Systems, Measurement and Control*, C94: 3–4, 1972.

[48] L.A. Zadeh. Outline of a new approach to the analysis of complex systems and decision processes. IEEE trans. Systems, Man, and Cybernetics, SMC-3: 28–44, 1973.

[49] L.A. Zadeh. The concept of a linguistic variable and its application to approximate reasoning (Part 1). *Information Sciences*, 8: 199–249, 1975.

[50] L.A. Zadeh. The concept of a linguistic variable and its application to approximate reasoning (Part 2). *Information Sciences*, 8: 301–357, 1975.

[51] L.A. Zadeh. The concept of a linguistic variable and its application to approximate reasoning (Part 3). *Information Sciences*, 9: 43–80, 1976.

[52] L.A. Zadeh. Fuzzy sets as a basis for a theory of possibility. *Fuzzy Sets and Systems*, 1:3–28, 1978.

[53] L.A. Zadeh. A theory of approximate reasoning. In D. Michie and L.I. Mikulich, editors, *Machine Intelligence 9*, New York: Halstead Press, 149–194, 1979.

[54] L.A. Zadeh. The role of fuzzy logic in the management of uncertainty in expert systems. *Fuzzy Sets and Systems*, 11:199–227, 1983.

Proceedings - International Conference on Fuzzy Logic & Neural Net...s (Iizuka, Japan, July 20-24, 1990) pp. 735—738

SIMILARITY-BASED INTERPRETATIONS OF FUZZY-LOGIC CONCEPTS

Enrique H. Ruspini
Artificial Intelligence Center
SRI International
Menlo Park, California, U.S.A.

INTRODUCTION

In this brief communication, we summarize the results of recent research on the conceptual foundations of fuzzy logic [5]. This research resulted in the formulation of several semantic models that interpret the major concepts and structures of fuzzy logic in terms of the more primitive notion of resemblance and similarity between "possible worlds," i.e., the possible states, situations or behaviors of a real-world system. The metric structures representing this notion of proximity are generalizations of the accessibility relation of modal logics [1].

Possibilistic reasoning methods may be characterized, by means of our interpretation, as approaches to the description of the relations of proximity that hold between possible system states that are logically consistent with existing evidence, and other situations, which are used as reference landmarks By contrast, probabilistic methods seek to quantify, by means of measures of set extension, the proportion of the set of possible worlds where a proposition is true.

Our discussion will focus primarily on the principal characteristics of a model, discussed in detail in a recent technical note [2], that quantifies resemblance between possible worlds by means of a similarity function that assigns a number between 0 and 1 to every pair of possible worlds. Introduction of such a function permits to interpret the major constructs and methods of fuzzy logic: conditional and unconditional possibility and necessity distributions and the generalized modus ponens of Zadeh on the basis of related metric relationships between subsets of possible worlds.

THE APPROXIMATE REASONING PROBLEM

Our semantic model of fuzzy logic is based on two major conceptual structures: the notion of possible world, which is the basis for our unified view of the approximate reasoning problem [3], and a metric structure that quantifies similarity between pairs of possible worlds.

If a reasoning problem is thought of as being concerned with the determination of the truth-value of a set of propositions that describe different aspects of the behavior of a system, then a *possible world* is simply a function (called a valuation) that assigns a unique truth value to every proposition in that set and that, in addition, is consistent with the rules of propositional logic. The set of all such possible worlds is called the *universe of discourse.*

In any reasoning problem, knowledge about the characteristics of the class of systems being studied combined with observations about the particular system under consideration restricts the extent of possible worlds that must be considered to a subset of the universe of discourse, called the *evidential set*, which will be denoted $\mathcal{E}$.

The purpose of the inferential procedures utilized in any reasoning problem may be characterized as that of establishing if, for a given proposition p, either $\mathcal{E} \Rightarrow p$ or $\mathcal{E} \Rightarrow \neg p$, i.e., whether existing evidence implies the hypothesis or its negation. In approximate reasoning problems, as illustrated in Figure 1, such determination is, by definition, impossible: there are some possible worlds in the the evidential set where the hypothesis is true and some where it is false.

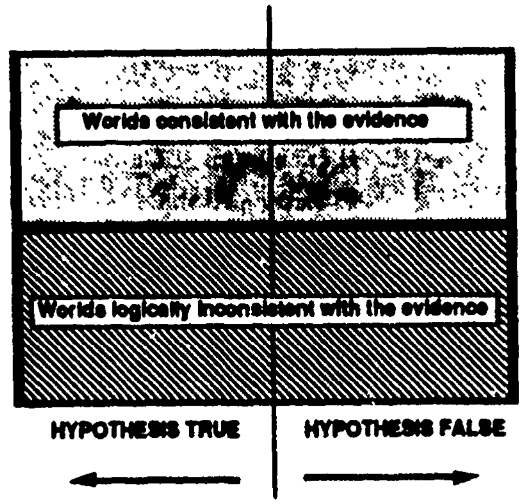

Figure 1: The approximate reasoning problem

SIMILARITY FUNCTIONS AND GENERALIZED IMPLICATION

In the view of fuzzy logic proposed by our model the purpose of possibilistic methods is the description of the evidential set by characterization of the resemblance relations that hold between its elements and elements of other sets used as reference landmarks.

To represent similarity or resemblance between possible worlds we introduce a binary function S that assigns a value between 0 and 1 to every pair of possible worlds w and w'. A value of S equal to 1 means that w and w' are identical while a value of S equal to 0 indicates that knowledge of propositions that are true in one possible world does not provide any indication about the nature of the propositions that are true in the other.

In addition to the above requirement of reflexivity, i.e. $S(w,w) = 1$, we will need to impose additional axioms to assure that S captures the semantics of a similarity relation. In addition to assuming that S is symmetric, i.e., $S(w,w') = S(w',w)$, we will also require that S satisfies a form of transitivity that is motivated by noting that if w, w' and w'' are possible worlds and if w is highly similar to w' and w' is highly similar to w'', then it would be surprising if w and w'' were highly dissimilar. This consideration indicates that knowledge of $S(w,w')$ and $S(w',w'')$ should provide a lower bound for values of $S(w,w'')$, as expressed by the inequality

$$S(w,w'') \geq S(w,w') \circledast S(w',w''),$$

where $\circledast$ is a binary operator used to represent a real function that produces the required bound. If reasonable requirements are imposed upon the function $\circledast$, it is easy to show that it has the properties of triangular norms: a class of functions that play a major role in multivalued logics [4].

The generalized transitivity property expressed by the above inequality may be easier to understand as a classical triangular inequality if it is noted that the function $\delta = 1 - S$ has the properties of a metric. When $\circledast$ is the Lukasïewicz norm $a \circledast b = \max(a + b - 1, 0)$, then the transitivity property of S is equivalent to the well-known triangular property of distance functions. If $\circledast$ corresponds to the Zadeh triangular norm $a \circledast b = \min(a,b)$, then δ may be shown to satisfy the more stringent ultrametric inequality.

The correspondence between propositions and subsets of possible worlds simplifies the interpretation of the classical rule of modus ponens as a rule of derivation based on the transitive property of set inclusion. If three propositions p, q and r are such that the set of possible worlds where p is true is a subset of the set of possible worlds where q is true, and if such set is itself a subset of the set of worlds where r is true, then the modus ponens simply states that the set of p-worlds is a subset of the set of r-worlds.

The conventional relation of set inclusion, based on the binary truth-value structure of classical logic, allows only to state that a set of possible worlds is a subset of another or that it is not. Introduction of a metric structure in the universe of discourse, however, permits the quantification of the degree by which a set is included into another. Every set of possible worlds, as illustrated in Figure 2, is a subset of some

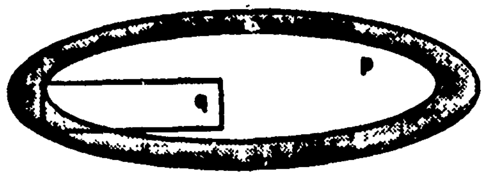

Figure 2: Degree of implication

neighborhood of any other set. The minimal amount of "stretching" that is required to include a set of possible worlds q in a neighborhood of a set of possible worlds p, given by the expression $\mathrm{I}(p\,|\,q) = \inf_{w' \vdash q} \sup_{w \vdash p} S(w, w')$, is called the *degree of implication*.

The degree of implication function has the important transitive property expressed by $\mathrm{I}(p\,|\,q) \geq \mathrm{I}(p\,|\,r) \circledast \mathrm{I}(r\,|\,q)$, which is the basis of the generalized modus ponens of Zadeh. As illustrated in Figure 3, this important rule of derivation tells us how much the set of p-worlds should be stretched to encompass q on the basis of knowledge of the sizes of the neighborhoods of p that includes r and of r that includes q.

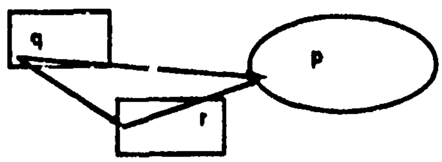

Figure 3: The generalized modus ponens

A notion dual to the degree of implication is that of *degree of consistence*, which quantifies the amount by which a set must be stretched to intersect another, and that is given by the expression $\mathrm{C}(p\,|\,q) = \sup_{w' \vdash q} \sup_{w \vdash p} S(w, w')$.

POSSIBILISTIC DISTRIBUTIONS

Although the transitive property of the degree of implication essentially provides the bases for the conceptual validity of the generalized modus ponens, this rule of derivation is typically expressed by means of necessity and possibility distributions.

An *unconditioned necessity distribution* given the evidence $\mathcal{E}$ is any function defined over propositions that bounds by below the degree of implication function, i.e., any function satisfying the inequality $\mathbf{Nec}(p) \leq \mathrm{I}(p\,|\,\mathcal{E})$. Correspondingly, an *unconditioned possibility distribution* is any upper bound for the degree of consistence function, i.e., $\mathbf{Poss}(p) \geq \mathrm{C}(p\,|\,\mathcal{E})$.

The definition of conditional possibility and necessity distributions makes use of a form of inverse of the triangular norm denoted $\oslash$ and defined by the expression

$$a \oslash b = \sup\{\, c :\ b \circledast c \leq a \,\}.$$

Using this function, it is possible to define conditional possibilistic distributions as follows:

Definition: A function $\mathbf{Nec}(\cdot|\cdot)$ is called a *conditional necessity distribution* for $\mathcal{E}$ if

$$\mathbf{Nec}(q|p) \leq \inf_{w \in \mathcal{E}} [\mathbf{I}(q\,|\,w) \oslash \mathbf{I}(p\,|\,w)]\,.,$$

Definition: A function $\mathbf{Poss}(\cdot|\cdot)$ is called a *conditional possibility distribution* for $\mathcal{E}$ if

$$\mathbf{Poss}(q|p) \geq \sup_{w \in \mathcal{E}} [\mathbf{I}(q\,|\,w) \oslash \mathbf{I}(p\,|\,w)]\,.$$

GENERALIZED MODUS PONENS

The *compositional rule of inference* or generalized modus ponens of of Zadeh is a generalization of the corresponding classical rule of inference that may be used even when known facts do not match the antecedent of a conditional rule. The interpretation provided by our model explains the generalized modus ponens as an extrapolation procedure that uses knowledge of the similarity between the evidence and a set of possible worlds p (the antecedent proposition), and of the proximity of p-worlds to q-worlds, to bound the similarity the latter to the evidential set. The actual statement of the generalized modus ponens for necessity distributions in terms of similarity structures makes use of a family $\mathcal{P}$ of satisfiable propositions that partitions the universe of discourse:

Theorem *(Generalized Modus Ponens for Possibility Functions)*: Let $\mathcal{P}$ be a partition and let q be a proposition. If $\mathbf{Poss}(p)$ and $\mathbf{Poss}(q|p)$ are real values, defined for every proposition p in $\mathcal{P}$, such that

$$\mathbf{Poss}(p) \geq \mathbf{C}(p\,|\,\mathcal{E})\,, \quad \mathbf{Poss}(q|p) \geq \sup_{w \in \mathcal{E}} [\mathbf{I}(q\,|\,w) \oslash \mathbf{I}(p\,|\,w)]\,,$$

then the following inequality is valid:

$$\sup_{\mathcal{P}} [\mathbf{Poss}(q|p) \circledast \mathbf{Poss}(p)] \geq \mathbf{C}(q\,|\,\mathcal{E})\,.$$

A dual result holds for necessity functions.

ACKNOWLEDGMENT

This work was supported in part by the US Air Force Office of Scientific Research under Contract No. F49620-89-K-0001 and in part by the US Army Research Office under Contract No. DAAL03-89-K-0156. The views, opinions and/or conclusions contained in this note are those of the author and should not be interpreted as representative of the official positions, decisions, or policies, either express or implied, of the Air Force Office of Scientific Research, the Department of the Army, or the United States Government.

REFERENCES

[1] G.E. Hughes and M.J. Creswell. *An Introduction to Modal Logic.* New York: Methuen, 1972.

[2] E.H. Ruspini. *On the Semantics of Fuzzy Logic.* Technical Note No. 475, Artificial Intelligence Center, SRI International, Menlo Park, California, 1989.

[3] E.H. Ruspini, *Approximate Reasoning: Past, Present, Future.* Information Sciences, forthcoming, 1990.

[4] E. Trillas and L. Valverde. On mode and implication in approximate reasoning. In M.M. Gupta, A. Kandel, W. Bandler, J.B. Kiszka, editors, *Approximate Reasoning and Expert Systems*, Amsterdam: North Holland, 157–166, 1985.

[5] L.A. Zadeh. A theory of approximate reasoning. In D. Michie and L.I. Mikulich, editors, *Machine Intelligence 9*, New York: Halstead Press, 149–194, 1979.

6th Workshop on Uncertainty in Artificial Intelligence, Boston, Massachusetts - U.S.A., July 1990

POSSIBILITY AS SIMILARITY: THE SEMANTICS OF FUZZY LOGIC

Enrique H. Ruspini
Artificial Intelligence Center
SRI International
Menlo Park, California, U.S.A.

Abstract

This paper addresses a number of fundamental issues on the nature of the concepts and structures of fuzzy logic, focusing, in particular, on the conceptual and functional differences that exist between probabilistic and possibilistic approaches.

A semantic model provides the basic framework allowing definition of possibilistic structures and concepts by means of a function that quantifies proximity, closeness, or resemblance between pairs of possible worlds. The resulting model is a natural extension, based on multiple conceivability relations, of the modal logic concepts of necessity and possibility. By contrast, typical, chance-oriented, probabilistic concepts and structures rely on measures of set extension that quantify the proportion of possible-worlds where a proposition is true.

Resemblance between possible worlds is quantified by a generalized similarity relation, i.e., a function that assigns a number between 0 and 1 to every pair of possible worlds. Using this similarity relation, which is a form of numerical complement of a classic metric or distance, the major constructs and methods of fuzzy logic—conditional and unconditional possibility and necessity distributions and the generalized modus ponens of Zadeh—are defined and interpreted.

This work was supported in part by the Air Force Office of Scientific Research under Contract No. F49620-89-K-0001 and in part by the United States Army Research Office under Contract No. DAAL03-89-K-0156.

1 Introduction

In this paper, we present a semantic model of the major concepts, structures, and methods of fuzzy or possibilistic [16,17] logic. This model is based on a framework that combines the notion of possible world [2] (i.e., a potential state or situation of a real-world system) with measures of proximity or resemblance between pairs of possible worlds. The resulting structures are substantially different in character and nature from those of probabilistic reasoning, which are based on measures of set extension, used to quantify the proportion of possible worlds where a given proposition is true.

The results reported in this paper are the latest in a continuing investigative effort aimed at clarifying basic conceptual similarities and differences between a numb of approaches to the treatment of imprecision and uncertainty. Using also possible-world semantic models, prior research has established that the Dempster-Shafer calculus of evidence may be interpreted by structures that result from the combination of conventional probabilistic calculus with epistemic logics [9]. By contrast, the formal structures discussed herein clearly show that fuzzy logic may be understood in a straightforward fashion using conventional metric notions in a space of possible worlds without resorting in any form to probabilistic concepts. Furthermore, the actual functions that are used to combine possibilistic knowledge are substantially different from those used in the probability calculus.

Our exposition, which will be limited to the major structures of fuzzy logic, defines possibilistic concepts using a more primive notion that has been found to be an essential component of important human cognitive processes [14]. The notion of *similarity*, in spite of its importance in reasoning proceses, has not received substantial attention in treatments based on the use of logical concepts.

Perhaps as a consequence of its reliance on metods for the manipulation of symbolic strings and on a single (partial order) relation between formulas (i.e., implication) as the basis for almost all of its techniques and procedures, there has been little attention given in formal logic to the consideration of other formal structures that capture important features of human knowledge such as the resemblance that exists between situations or circumstances. Although, for example, stating that Mary is worth $1,000,000 as oppossed to saying that she is worth $999,999 may be rather inconsequential in terms of the implication of either statement to a decision-maker (e.g., trying to establish a credit line), there is nothing in the basic framework of logic that makes the second statement any more different than saying that Mary is broke (i.e., neither of the three statements is logically consistent with the other two).

The determination and use of similarity information is, however, not only central to all forms of analogical reasoning but it is an essential element in the derivation of physical law. Formal studies in measurement theory [7] clearly show the role that measures based on similar behavior play in the derivation of rational measurement schemes, while also explaining the ubiquitous presence of numeric scales throughout science.

The results presented in this paper show that, when such notions of proximity are formalized in the context of a possible-worlds model, the major functional structures of fuzzy logic—possibility and necessity distributions—and its major inferential procedure—the generalized modus ponens of Zadeh—may be readily explained as a natural extension of classical logical concepts. In particular, possibility and necessity distributions simply correspond to best and worst scenarios in a space of possible real-world states, while the generalized modus ponens [17] is a sound inferential procedure that may be regarded as a form of logical extrapolation between neighboring situations.

The scope of this paper prevents a detailed discussion of all pertinent results and derivations. A complete account of all relevant matters regarding the similarity-based model of fuzzy logic presented in this paper is presented in a related technical note [10], which, essentially, this paper summarizes.

2 The Approximate Reasoning Problem

Our model of the approximate reasoning problem is based on the notion of "possible world." Informally, possible worlds are the conceivable states of affairs of a real-world system that are consistent with the laws of logic.

Restricting ourselves, for the sake of simplicity to propositional formulations, a possible world is a function [2] that assigns a unique conventional truth value (i.e., **true** or **false**) to every proposition that describes some relevant aspect of the state of the system and, that, in addition satisfies the axioms of propositional logic.

In the absence of any knowledge about the behavior of a system of interest or of any observation about its state, it is impossible to determine which, among all conceivable situations, corresponds to the actual state of the real world. Availability of factual evidence or determination of the laws of behavior of the system permits, however, to eliminate some possible worlds in this *universe of discourse* from consideration. The remaining possible worlds correspond to satisfiable propositions that, in addition, are logically consistent with the evidence. This subset of conceivable situations or scenarios will be called the *evidential set*, denoted $\mathscr{E}$.

If the typical reasoning problem is thought of as the determination of the truth value of a proposition h (the *hypothesis*), then an approximate reasoning problem may be described as one where available evidence does not permit such evaluation without ambiguity. In other words, as illustrated in Figure 1, there are some members of the evidential set where the hypothesis is true and some where it is false.

Our approach to the formalization of the major concepts and structures of fuzzy logic of fuzzy logic is based on a generalization of a central concept of semantic models of modal logics. Modal logics [4] may be generally described as extensions of conventional two-valued logic that permit to qualify, in various ways, the meaning of propositional truth.

In our model, we utilize modal concepts to explain basic possibilistic structures using the more primitive notion of similarity. This notion is introduced, however, by means of conventional set-theoretic and logical concepts. In this regard, our approach to the study

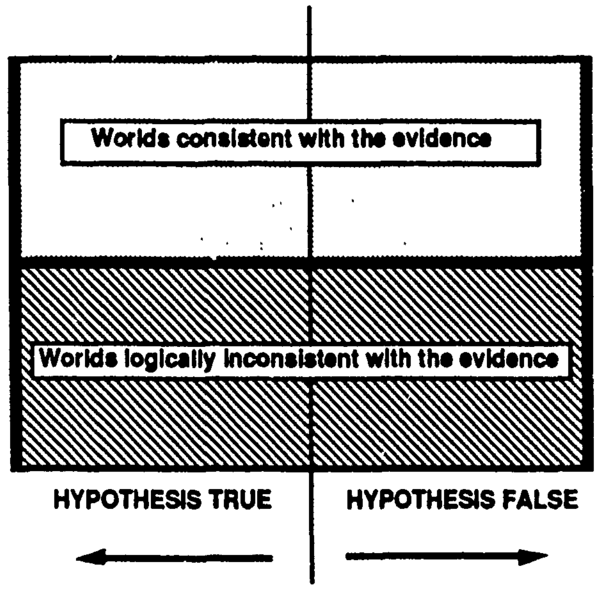

Figure 1: The Approximate Reasoning Problem.

of the interplay of modal and possibilistic logics is different from approaches such as that used by Lakoff [6] who sought to generalize modal logics using fuzzy-set concepts; or that of Dubois and Prade [3], who investigated modal structures with a view to the development of formal proof mechanisms in possibilistic logic.

A major concept of semantic models of modal logic systems is a binary relation R, called the *accessibility* or *conceivability* relation. This relation is assumed to have a number of properties intended to capture the semantics of various qualifications of propositional truth, ranging from logical necessity through the state of knowledge of rational agents to concepts related to the ideal behavior of ethical decision-makers.

Our aim is to characterize the extent by which statements that are true in one situation or scenario may be said, perhaps with some suitable modification, to be true in another state of affairs that resembles it. We are particularly interested in describing more general (i.e., less specific) propositions that are true in one possible world as a function of the propositions that are true in another. In order to model a continuous range of proximity between possible worlds, we will generalize the notion of accessibility relation to a full family of binary relations R_α, indexed by a numerical parameter α taking values between 0 and 1, along the same lines—albeit with a different purpose—utilized by Lewis in his treatment of counterfactuals [5].

3 Similarity and Graded Possibility

We will introduce a family of accessibility relations

$$\{R_\alpha : \alpha \in [0,1]\},$$

by means of a binary function S, called the *similarity relation*, that maps pairs of possible worlds into numbers between 0 and 1. The multiple relations of accessibility R_α are defined in terms of this similarity function by

$$wR_\alpha w' \text{ if and only if } S(w,w') \geq \alpha \qquad \alpha \in [0,1].$$

The function S is intended to capture a notion of proximity, closeness, or resemblance between possible worlds with a value of 1 corresponding to the identity of possible worlds and a value of 0 indicating that knowledge of propositions that are true in a possible world does not provide any indication of the propositions that are true in the other. To assure that the function S has the semantics of a relation that quantifies resemblance between possible states of affairs, it is necessary to require that it satisfies a number of properties.

Besides the above mentioned property that the similarity between a possible world and itself has the highest possible value, equivalent to stating that each accessibility relation R_α is *reflexive*, we will also require that the similarity between *different* possible worlds be strictly less than one. This requirement is intended to assure that the similarity relation may distinguish between different states of the possible world.

The similarity relation will also be assumed to be *symmetric*, and to satisfy a relaxed form of transitivity. Clearly, if the pairs of possible worlds (w,w') and (w',w'') correspond to highly similar situations, it would be surprising if w and w'' were highly dissimilar. It is natural to assume, therefore, that

$$S(w,w'') \leq S(w,w') \circledast S(w',w''),$$

where $\circledast$ is a binary operator used to represent the lower bound as a function of its arguments. This requirement is equivalent to the relaxed transitivity condition

$$R_{\alpha \circledast \beta} \subseteq R_\alpha \circ R_\beta,$$

which replaces the usual, more stringent, definition of transitivity.

Imposition of reasonable requirements upon the function $\circledast$ shows that it has the properties of a *triangular norm* [11]. These functions, which play a significant role in multivalued logics [12], may be justified, therefore, purely on the basis of metric considerations. Important examples of triangular norms are

$$a \circledast b = \min(a,b), \qquad a \circledast b = \max(a+b-1,0), \qquad \text{and} \qquad a \circledast b = ab,$$

called the *Zadeh*, *Lukasïewicz*, and *product* triangular norms, respectively.

The generalized transitivity property that is expressed by triangular norms clarifies their relationship to the conventional mathematical concept of metric. If S is a similarity function, then the function $\delta = 1 - S$ has the properties of a distance function. When $\circledast$ corresponds to the Lukasïewicz norm, then the transitivity property of S corresponds to the well-known triangular property of distance functions. If $\circledast$ correponds to the Zadeh triangular norm, then δ may be shown to satisfy the more stringent *ultrametric* inequality.

4 Degree of Implication and Consistence

5 Variables

6 Possibility and Necessity Distributions

7 The Generalized Modus Ponens

8 Conclusions

References

[1] J.F. Baldwin. A new approach to approximate reasoning using a fuzzy logic. *Fuzzy Sets and Systems*, 2:302–335, 1979.

[2] R. Carnap. *The Logical Foundations of Probability.* Chicago: University of Chicago Press, 1950.

[3] D. Dubois and H. Prade. In search of a modal system for possibility theory. In *Proc. 8th. European Conf. on Artificial Intelligence*, 501–506, Munich: Technical University, 1988.

[4] G.E. Hughes and M.J. Creswell. *An Introduction to Modal Logic.* New York: Methuen, 1972.

[5] D. Lewis. *Counterfactuals.* Cambridge, Massachusetts: Harvard University Press, 1973.

[6] G. Lakoff. Hedges: A study in meaning criteria and the logic of fuzzy concepts. *J. Philos. Logic*, 2: 458–508, 1973.

[7] Narens, L. and R. Duncan Luce, *Measurement: The Theory of Numerical Assignments*, Psych. Bull., 99, 166–180, 1986.

[8] Z. Pavlak. Rough sets. *Int. J. Comput. Inf. Sci.*, 11:341–356, 1982.

[9] E.H. Ruspini. *The Logical Foundations of Evidential Reasoning.* Technical Note No. 408, Artificial Intelligence Center, SRI International, Menlo Park, California, 1987.

[10] E.H. Ruspini. *On the Semantics of Fuzzy Logic.* Technical Note No. 475, Artificial Intelligence Center, SRI International, Menlo Park, California, 1989.

[11] B. Schweizer and A. Sklar. Associative functions and abstract semigroups. *Publ. Math. Debrecen*, 10:69–81, 1963.

[12] E. Trillas and L. Valverde. On mode and implication in approximate reasoning. In M.M. Gupta, A. Kandel, W. Bandler, J.B. Kiszka, editors, *Approximate Reasoning and Expert Systems*, Amsterdam: North Holland, 157–166, 1985.

[13] L. Valverde. On the structure of F-indistinguishability operators. *Fuzzy Sets and Systems*, 17:313–328, 1985.

[14] A. Tversky. Features of similarity. *Psychological Review*, 84:433–460, 1977.

[15] L.A. Zadeh. Fuzzy sets. *Information and Control*, 8: 338–353, 1965.

[16] L.A. Zadeh. Fuzzy sets as a basis for a theory of possibility. *Fuzzy Sets and Systems*, 1:3–28, 1978.

[17] L.A. Zadeh. A theory of approximate reasoning. In D. Michie and L.I. Mikulich, editors, *Machine Intelligence 9*, New York: Halstead Press, 149–194, 1979.

Third International Conference on Information
Processing and Management of Uncertainty by
Expert Systems

SIMILARITY MODELS FOR FUZZY LOGIC

Enrique H. RUSPINI

Artificial Intelligence Center
SRI International
Menlo Park, California - U.S.A.

1 Introduction

The notion of similarity, which plays a major role in human cognitive processes [4], may be used to formulate a number of semantic models that explain the major concepts of fuzzy logic [5]. These formalisms show that possibilistic reasoning is fundamentally different from approaches based on the notion of probability: an additive measure of set extent.

The idea that knowledge of propositions that are true in certain situations may be used to derive truth-values in similar situations has not received much attention in conventional logical treatments. This state of affairs may be traced to the reliance of logical methods on symbolic procedures that only recognize one important relationship between formulas, i.e., the partial order defined by implication.

This paper briefly describes one model t' t explains possibilistic structures: possibility and necessity distributions; and the major derivational rule of fuzzy logic: the generalized modus ponens; in terms of simpler concepts related to notions of resemblance between possible worlds. In particular, the latter procedure is shown to generalize its classical counterpart by allowing a form of logical extrapolation between similar situations and scenarios. A full discussion of this model and its implications is presented in a related technical note [2].

2 The Approximate Reasoning Problem

Our model is based on a unified view of approximate reasoning methodologies that regards these procedures as techniques that describe certain properties of subsets of possible worlds. Informally, possible worlds are the conceivable states (i.e., scenarios, situations) of a real-world system that are consistent with the laws of logic. Restricting ourselves to propositional formulations, a possible world is a function that assigns a unique truth value (i.e., **true** or **false**) to every proposition that describes a relevant aspect of system of state and behavior and that, in addition, satisfies the axioms of propositional logic.

The set of all such possible worlds is called the *universe of discourse*. Knowledge about the class of systems being studied, combined with observations about the actual system under consideration, usually restricts the set of states that must be considered in an approximate reasoning problem to a proper subset of this universe. This subset, denoted $\mathcal{E}$, is called the ***evidential set***.

In a typical approximate reasoning problem, as illustrated in Figure 1, available evidence does not permit to determine if a hypothesis of interest is true or false. Being unable to determine such truth value, approximate reasoning methods try to describe significant properties of the evidential set. Possibilistic techniques describe the relations of similarity that hold between possible worlds in the evidential set and possible worlds in other sets, used as reference landmarks.

3 Similarity Relations

To capture the notion of proximity or resemblance between possible worlds, we will introduce a function S that assigns a number between 0 and 1 to every pair of possible worlds. This function permits to define a family of relations between possible worlds that generalizes the classical modal notion of *accessibility* [1]. By assumption, S attains a value of 1 only when its two arguments are identical. A value of 0, by contrast, is intended to

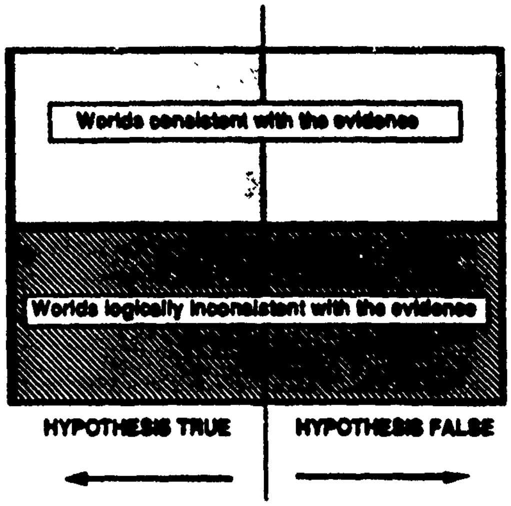

Figure 1: The approximate reasoning problem

model the fact that knowledge of propositions that are true in one possible world does not provide any indication about propositions that are true in the other.

The similarity relation will also be assumed to be symmetric and to satisfy a relaxed form of transitivity, intended to capture the notion that the similarity between two possible worlds w and w'' bears some relation to the values of the similarities between each of them and a third world w', expressed by the inequality

$$S(w, w'') \geq S(w, w') \circledast S(w', w''),$$

where $\circledast$ is a binary operator defined for pairs of numbers in $[0, 1]$. Imposition of reasonable requirements upon the function $\circledast$ shows that it has the properties of a *triangular norm* [3].

In what follows, we will also need a form of inverse of the triangular norm $\circledast$, denoted $\oslash$, and defined by the expression

$$a \oslash b = \sup\{ c : b \circledast c \leq a \}.$$

4 Degree of Implication and Degree of Consistence

The classical rule of modus ponens may be thought of as expressing the transitive property of subset inclusion. Introducing a metric relation and its associated topology permits to extend this relation by measures that quantify the size of the neighborhood of a set that contains another set. We will say that q implies p to the degree α if, for every q-world w there exists a p-world w' such that $S(w, w') \geq \alpha$. Since it is true that $S(w, w') \geq 0$ for every pair of possible worlds, it is obvious that any proposition implies any other proposition to some degree.

Informally, the definition of graded implication means that if p is stretched to the degree α, then this stretched set will include q. The upper bound of the values α such that q implies p to the degree α, expressed by

$$\mathbf{I}(p \mid q) = \inf_{w \vdash q} \sup_{w' \vdash p} S(w, w'),$$

defines a function $\mathbf{I}$ called the *degree of implication*. The degree of implication, which is related to the notion of Hausdorff distance, has the transitive property

$$\mathbf{I}(p \mid q) \geq \mathbf{I}(p \mid r) \circledast \mathbf{I}(r \mid q),$$

which is the basis of the *generalized modus ponens* of Zadeh, illustrated in Figure 2.

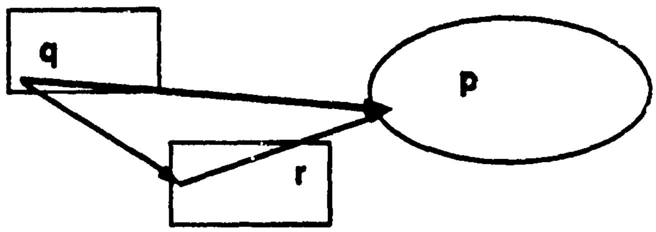

Figure 2: The generalized modus ponens

A notion that is dual to that of the degree of implication is the *degree of consistence*, which quantifies the amount by which a set must be stretched in order to intersect another set,

$$\mathbf{C}(p \mid q) = \sup_{w \vdash q} \sup_{w' \vdash p} S(w, w').$$

Obviously,

$$\mathbf{I}(p \mid q) \leq \mathbf{C}(p \mid q).$$

5 Possibility and Necessity Distributions

An *unconditioned necessity* distribution for $\mathcal{E}$ is any function $\mathbf{Nec}(\cdot)$ mapping propositions (i.e., subset of possible worlds) into numbers between 0 and 1, such that

$$\mathbf{Nec}(p) \leq \mathbf{I}(p \mid \mathcal{E}),$$

i.e., a lower bound of the degree of implication of p by $\mathcal{E}$. Correspondingly, an *unconditioned possibility* distribution is an upper bound for the degree of consistence of p and $\mathcal{E}$, i.e.,

$$\mathbf{Poss}(p) \geq \mathbf{C}(p \mid \mathcal{E}).$$

Unconditioned necessity and possibility distributions measure how much a set must be stretched to enclose or intersect, respectively, the evidential set. The conditional counterparts of these notions characterize the proximity relations that exist between evidential worlds and worlds satisfying a consequent proposition q as a proportion of the similarity that exists between those evidential worlds and worlds that satisfy the antecedent proposition p.

A function $\mathbf{Nec}(\cdot|\cdot)$ is called a *conditional necessity distribution* for $\mathcal{E}$ if

$$\mathbf{Nec}(q|p) \leq \inf_{w \vdash \mathcal{E}} [\mathbf{I}(q\,|\,w) \oslash \mathbf{I}(p\,|\,w)].$$

Correspondingly, a function $\mathbf{Poss}(\cdot|\cdot)$ is called a *conditional possibility distribution* for $\mathcal{E}$ if

$$\mathbf{Poss}(q|p) \geq \sup_{w \vdash \mathcal{E}} [\mathbf{I}(q\,|\,w) \oslash \mathbf{I}(p\,|\,w)].$$

6 The Generalized Modus Ponens

The usual statement of the *compositional rule of inference* or generalized modus ponens of of Zadeh [5] is made in terms of a relationship between unconditioned and conditioned distributions rather than in its simpler form, given above, as the transitive property of the degree of implication.

The generalized modus ponens is a sound logical extrapolation procedure that uses information about the metric relations that hold between different subsets. On the basis of information about the similarity between evidential worlds and a set of possible worlds (i.e., the antecedent proposition p), and of knowledge about the relative proximity of p-worlds and q-worlds (i.e., conditional distributions), the generalized modus ponens produces bounds for the similarity between evidential worlds and those that satisfy the consequent proposition q (i.e., unconditioned distributions for the consequent).

The actual formal statement of the generalized modus ponens makes use of the notion of *partition* of the universe of discourse. A partition $\mathcal{P}$ simply corresponds to an ordinary partition of of the universe of discourse into disjoint subsets, or, equivalently to a collection of mutually disjoint propositions such that their disjunction is always true.

Using this concept, the generalized modus ponens for possibility distributions may be stated as follows in terms of distributions defined using similarity structures:

Theorem: Let $\mathcal{P}$ be a partition and let q be a proposition. If $\mathbf{Poss}(p)$ and $\mathbf{Poss}(q|p)$ are real values, defined for every proposition p in $\mathcal{P}$, such that

$$\begin{aligned} \mathbf{Poss}(p) &\geq \mathbf{C}(p\,|\,\mathcal{E}), \\ \mathbf{Poss}(q|p) &\geq \sup_{w \vdash \mathcal{E}} [\mathbf{I}(q\,|\,w) \oslash \mathbf{I}(p\,|\,w)], \end{aligned}$$

then the following inequality is valid

$$\sup_{\mathcal{P}} [\mathbf{Poss}(q|p) \circledast \mathbf{Poss}(p)] \geq \mathbf{C}(q\,|\,\mathcal{E}).$$

A dual result holds for necessity distributions.

7 Conclusion

Similarity models provide useful interpretations for the basic concepts of possibilistic logic using a more primitive notion than that of possibility. In addition to clearly showing that fuzzy logic structures are not related to probabilistic notions, the resulting framework provides a solid basis for the study and extension of possibilistic logic in a number of directions of considerable practical importance.

8 Acknowledgment

This work was supported in part by the Air Force Office of Scientific Research under Contract No. F49620-89-K-0001 and in part by the United States Army Research Office under Contract No. DAAL03-89-K-0156.

The views, opinions and/or conclusions contained in this paper are those of the author and should not be interpreted as representative of the official positions, decisions, or policies, either express or implied, of the Air Force Office of Scientific Research, the Department of the Army, or the United States Government.

References

[1] G.E. Hughes and M.J. Creswell. *An Introduction to Modal Logic*. New York: Methuen, 1972.

[2] E.H. Ruspini. *On the Semantics of Fuzzy Logic*. Technical Note No. 475, Artificial Intelligence Center, SRI International, Menlo Park, California, 1989.

[3] E. Trillas and L. Valverde. On mode and implication in approximate reasoning. In M.M. Gupta, A. Kandel, W. Bandler, J.B. Kiszka, editors, *Approximate Reasoning and Expert Systems*, Amsterdam: North Holland, 157–166, 1985.

[4] A. Tversky. Features of similarity. *Psychological Review*, 84:433–460, 1977.

[5] L.A. Zadeh. Fuzzy sets as a basis for a theory of possibility. *Fuzzy Sets and Systems*, 1:3–28, 1978.

Third International Conference on Information
Processing and Management of Uncertainty by
Expert Systems

ON EVIDENTIAL REASONING

John D. LOWRANCE, Enrique H. RUSPINI, Thomas M. STRAT

Artificial Intelligence Center
SRI International
Menlo Park, California - U.S.A.

1 Introduction

If artificially intelligent systems are to produce adequate assessments of the state and behavior of the real world, they must cope with information and knowledge that is characterized by varying degrees of uncertainty, ignorance, and correctness. To address this need, we have developed a technology called *evidential reasoning*. It is formally based upon the Dempster-Shafer theory of belief functions; it has been implemented as a domain-independent automated reasoning system; it has been successfully applied to a range of real-world problems [2]. Yet, its reliance on belief functions has drawn criticism.

Our choice of an approach based on the Dempster-Shafer theory was not arbitrary. We believe that it has important methodological advantages such as its ability to represent ignorance in a direct and straightforward fashion, its consistency with classical probability theory, its compatibility with Boolean logic, and its manageable computational complexity. At the same time, we recognize that other approaches may also complement and augment the assessments provided by evidential reasoning.

We will examine, within the limited scope provided by the format of this paper, several criticisms of belief functions that have appeared in the literature. We plan, however, a more thorough discussion of these criticisms in a related volume to be published in connection with this conference.

We discuss first the fundamental theoretical bases supporting the belief-function approach and justify its use in terms of the requirements imposed by ignorance of certain probability distributions. We consider the nature of Dempster's rule of combination and argue that negative assessments either misinterpret the nature of the distributions being combined or ignore the basic independence assumptions that assure its validity.

We answer also to critiques based on the computational complexity of the belief-function approach. Such criticisms claim that the complexity of probabilistic knowledge representations grows exponentially with the size of the frame, thus making the theory unsuited for automated reasoning. Other comments addressed in our presentation center on limitations on the representational ability of belief functions and the lack of certain methodological capabilities (e.g., decision-making mechanisms).

Despite the criticism that belief functions have drawn, we believe evidential reasoning to be well-founded and to have practical utility in a broad range of applications.

2 On Theoretical Soundness

The theory of belief functions was originated by Dempster [1] in the context of statistical research. The use of the term "belief," together with its subjectivist connotations, is due to Shafer [7], who first applied the theory to the analysis of the information contained in imprecise and uncertain evidence.

Although much skepticism has been voiced about the naturality of belief functions and their agreement with conventional probabilistic approaches, its theoretical bases are provided by a simple consideration about the role of evidence as a basic information carrier.

In classical probabilistic treatments, it is assumed that, under certain evidential conditions $\mathscr{E}$, the value $\Pr(p|\mathscr{E})$ of the likelihood of a particular statement p is known. This view of evidence, while adequate to represent the informational conditions of most controlled experimental setups, fails, however, to adequately model the effects that acquiring similar information has on our state of knowledge when the state of the world could not be so readily controlled.

In such circumstances, whenever the evidence $\mathscr{E}$ is observed, three possible informational outcomes may result from examination of further information that later turns out to improve our state of knowledge: either p is found to be true, $\neg p$ is found to be true (i.e., p is false), or such information is insufficient to determine the truth value of p. Use of modal logic concepts, which are the bases of the formal model of Ruspini [6], suggests the use of the notation $\mathbf{K}p$, $\mathbf{K}\neg p$, and $\mathbf{I}p$ to identify these outcomes. Since these alternatives are exclusive, it is clear that

$$\Pr(\mathbf{K}p) + \Pr(\mathbf{K}\neg p) + \Pr(\mathbf{I}p) = 1.$$

As shown by Ruspini, the function $\mathrm{Bel}(p) = \Pr(\mathbf{K}p)$, has the properties of a belief function, as axiomatized by Shafer. Furthermore, since it is possible that $\Pr(\mathbf{I}p) > 0$, then,in general, it is $\mathrm{Bel}(p) + \mathrm{Bel}(\neg p) \leq 1$. This inequality follows naturally, therefore, from classical probability theory, applied here to considerations about the *provability* of certain propositions, as called by Pearl [4].

Similar considerations about the informational effect of independent bodies of evidence, which are beyond the scope of this short summary, indicate that Dempster's combination formula is, under its stated assumptions, completely consistent with conventional probability calculus.

This interpretation quickly disposes of erroneous arguments based on unintended interpretations of the intervals defined by belief functions. Each such interval represents ignorance of a single probability value for a fixed proposition p under fixed evidential conditions $\mathscr{E}$. If critics choose to interpret such intervals as the possible values that conditional probalilities might attain when further evidence is collected, as suggested by Pearl [3], belief functions will not, indeed, behave according to such unintended semantics.

3 On Decision Support

A criticism of a more fundamental nature, however, is often raised regarding the epistemological need for the belief-function approach. Summarized by statements such as Pearl's [4] question: "why we should concern ourselves with the probability that the evidence implies A, rather than the probability that A is true, given the evidence?," these arguments correctly point to the basic knowledge requirement that most decision problems entail: if a rational choice is to be made, then we must have a proper informational basis to do it.

This obvious consideration is twisted, however, to argue for the necessity to estimate unknown probability values when they are not available. We do not think that this modified, or *pragmatic necessity*, argument is either sound or compelling. To answer Pearl's question, we concern ourselves with the probability of provability because that is all that our data and the laws of logic can provide. We would rather measure the probabilities of truth, and endeavor to do so whenever possible, but we do not think, however, that probabilities should be guessed, simply because we are compelled to choose a course of action, anymore than any other unknown physical parameter value.

In our view, the belief-function approach may be used in a straightforward fashion to produce intervals of possible utility values. When such intervals overlap and cannot be ordered, this fact simply reflects a basic defficiency in our knowledge. We look down upon "pragmatic justifications" with the same concern that any experimental scientist shows about proposals to guess what he has not measured: the ability to make decisions in the absence of knowledge is, in our view, a handicap rather than an advantage of a method.

4 On the Dempster Formula

The Dempster formula is, currently, the principal evidence integration mechanism of the belief-function approach. It was derived in the context of a basic model of the effect of probabilistic evidence that correctly interprets such evidence as constraints on probability values rather than as the source of the actual values, which are typically undetermined.

The formula may be described as an expression that yields bounds for the conditional probability distribution $\Pr(\cdot|\mathscr{E}_1, \mathscr{E}_2)$ on the basis of similar bounds for the probability distributions $\Pr(\cdot|\mathscr{E}_1)$ and $\Pr(\cdot|\mathscr{E}_2)$, under certain conditions of independence.

Criticisms about the Dempster formula may be broadly characterized as being the consequence of two basic misunderstandings about its validity and generality.

First, the formula is intended to be applied only to those situations where its underlying assumptions are valid. Alleged counterexamples such as that of the "three prisoner problem," referenced by Pearl [4], fail to satisfy such assumptions and cannot be correctly said to be theoretical failures. We agree with Pearl, however, in its criticism of the use of the Dempster formula to produce a *conditioning formula*, leading to counterintuitive results (the "spoiled sandwich" effect), which we consider also to reflect failure of the basic independence assumptions. We are endeavoring, however, to extend the original theory to produce expressions to produce and utilize conditional belief information [5].

The second type of criticisms are based on the erroneous assumption that the two evidential bodies being combined should be interpreted as bounds, provided by two independent "experts," which constrain the values of the same probability distribution. As it was pointed out before, the formula combines two different conditional probability distributions.

5 On Generality and Complexity

The lack of generality of the belief-function approach to represent interval constraints on a family of probability distributions is well known. Our reliance on the belief-function approach, in spite of such lack of generality, is based on two major considerations.

First, our experience shows that, notwithstanding criticisms based on unrealistic worst-case scenarios, the approach is computationally efficient. In particular, we have found that representation of belief functions in terms of basic probabilistic assignments results in a storage and manipulation scheme that is both economical and easy to understand. In addition, we have sucessfully implemented tools, such as summarization and coarsening operators, which may be effectively utilized to limit representational complexity.

Second, our current functional operators have been chosen to guarantee that probabilistic information will always be capable of being represented within the scope of the approach, as more general constraints do not either enter into consideration or appear as the result of any of its functions.

Our current concerns with the manipulation of conditional and dependent evidence show, however, that, for some important problems, the results of evidential combination fall outside the scope of its representational capabilities. Although more general schemes, such as interval probabilities, do not suffer from this limitation, their inherent complexity precludes their practical application.

Ongoing research indicates, on the other hand, that the belief-function approach may be used to approximate the results of these general evidential combination operations. This research also shows the basic errors inherent in criticisms that regard the belief-function approach as a fully developed methodology incapable of sustaining further enhancement and modification. Having been studied in depth for only fifteen years, its technological status is that of a young discipline being capable of enhancement on its own and of combination with other approaches to produce more general tools for probabilistic reasoning. Far from proving that we have reached a technological plateau, our investigations indicate that much is yet to be gained from such a development and integration process.

6 Acknowledgment

This work was supported in part by the Defense Advanced Research Project Agency under contract N00039-88-C-0248, in part by the Air Force Office of Scientific Research under Contract No. F49620-89-K-0001, and in part by the United States Army Research Office under Contract No. DAAL03-89-K-0156. The views, opinions and/or conclusions contained in this paper are those of the authors and should not be interpreted as representative of the official positions, decisions, or policies, either express or implied, of the supporting agencies or the United States Government.

References

[1] A.P. Dempster. Upper and lower probabilities induced by a multivalued mapping. *Annals of Statistics*, 38:325–339, 1967.

[2] J. Lowrance, T. Strat and T. Garvey. A Framework for Evidential-Reasoning Systems. *Proc. National Conference on Artificial Intelligence*, 1986, pp. 896–903.

[3] J. Pearl. On probability intervals. *International Journal of Approximate Reasoning*, 2:211–216, 1988.

[4] J. Pearl. Reasoning with Belief Functions: An Analysis of Compatibility. *International Journal of Approximate Reasoning*, forthcoming.

[5] E.H. Ruspini. Approximate Deduction in Single Evidential Bodies. *Proceedings 1986 AAAI Uncertainty Workshop*, Philadelphia, 1986.

[6] E.H. Ruspini. *The Logical Foundations of Evidential Reasoning*. Technical Note No. 408, Artificial Intelligence Center, SRI International, Menlo Park, California, 1987.

[7] G. Shafer. *A Mathematical Theory of Evidence*. Princeton University Press, 1976.

UNDERSTANDING EVIDENTIAL REASONING

John D. Lowrance
Enrique H. Ruspini
Thomas M. Strat

Artificial Intelligence Center
SRI International
Menlo Park, California, U.S.A.

Abstract

We address recent criticisms of evidential reasoning: an approach to the analysis of imprecise and uncertain information that is based on the Dempster-Shafer calculus of evidence.

We show that evidential reasoning can be interpreted in terms of classical probability theory and that the Dempster-Shafer calculus of evidence may be considered to be a form of generalized probabilistic reasoning based on the representation of probabilistic ignorance by intervals of possible values. In particular, we emphasize that it is not necessary to resort to nonprobabilistic or subjectivist explanations to justify the validity of the approach.

We answer to conceptual criticisms of evidential reasoning primarily on the basis of their confusion between the current state of development of the theory — mainly theoretical limitations in the treatment of conditional information— with its potential usefulness to treat a wide variety of uncertainty-analysis problems. Similarly, we indicate that the supposed lack of decision-support schemes of generalized probability approaches is not a theoretical handicap but, rather, an indication of basic informational shortcomings that is a desirable asset of any formal approximate reasoning approach. We also point to potential shortcomings of the underlying representation scheme to treat general probabilistic reasoning problems.

We consider also methodological criticisms of the approach focusing primarily on the alleged counterintuitive nature of Dempster's combination formula showing that such results are the result of its misapplication. We address also issues of complexity and validity of scope of the calculus of evidence.

1 Introduction

If artificially intelligent systems are to produce adequate assessments of the state and behavior of the real world, they must cope with information and knowledge that is characterized by varying degrees of uncertainty, ignorance, and correctness. To address this need, we have developed a technology called *evidential reasoning*. It is formally based upon the Dempster-Shafer theory of belief functions; it has been implemented as a domain-independent automated reasoning system; it has been successfully applied to a range of real-world problems [11]. Yet, its reliance on belief functions has drawn criticism.

Our choice of an approach based on the Dempster-Shafer theory was not arbitrary. We believe that it has important methodological advantages such as its ability to represent ignorance in a direct and straightforward fashion, its consistency with classical probability theory, its compatibility with Boolean logic, and its manageable computational complexity. At the same time, we recognize that other approaches may also complement and augment the assessments provided by evidential reasoning.

We examine several criticisms of belief functions that have appeared in the literature, discussing first the fundamental theoretical bases supporting the belief-function approach and justifying its use in terms of the requirements imposed by ignorance of certain probability distributions. We consider the nature of Dempster's rule of combination and argue that negative assessments either misinterpret the nature of the distributions being combined or ignore the basic independence assumptions that assure its validity. We stress also that it is not necessary to rely on explanations that are either nonprobabilistic or subjective to justify the validity of the Dempster-Shafer calculus of evidence.

Furthermore, we show that certain apparently counterintuitive properties of the approach (e.g., the "spoiled sandwich" paradox) are the natural consequence of considering families of possible probability distributions that solve an approximate reasoning problem. In the context of this discussion, we indicate also the inherent pitfalls of "axiomatic" approaches that accept or reject methodologies on the basis of their compliance with allegedly intuitive principles.

We answer also to critiques based on the computational complexity of the belief-function approach. Such criticisms claim that the complexity of probabilistic knowledge representations grows exponentially with the size of the frame, thus making the theory unsuited for automated reasoning. Other comments addressed in our presentation center on limitations on the representational ability of belief functions and the lack of certain methodological capabilities (e.g., decision-making mechanisms).

Despite the criticism that belief functions have drawn, we believe that evidential reasoning is well-founded and that it may be effectively applied to the solution of a broad range of important practical problems.

Most of our comments will be made in direct reply to a recent criticism of the belief-function approach by Pearl [15] since we feel that his paper encompasses most of the major worries and concerns with the calculus of evidence. While most of the discussion of this paper consists of direct responses to issues raised by Pearl and others, our overall objective is considerably broader. Our answers are motivated by the same remarks of DeGroot, quoted by Pearl at the conclusion of his work, about the need to use our methodological approaches "... with the utmost care and in accordance with the highest ethical standards." Our aim, like Pearl's, is to enlighten and clarify, through careful discussion of rather subtle and delicate issues, rather than to engage in dogmatic defense of one approach to the detriment of another. It is our earnest hope that this work, in conjunction with other evaluations of the belief-function approach, will help to understand its bases, capabilities, and limitations.

2 On Theoretical Soundness

The theory of belief functions was originated by Dempster [4] in the context of statistical research. The use of the term "belief," together with its subjectivist connotations, is due to Shafer [18], who first applied the theory to the analysis of imprecise and uncertain evidence.

Although much skepticism has been voiced about the naturality of belief functions and their agreement with conventional probabilistic approaches, its theoretical bases are provided by a simple consideration about the role of evidence as a basic information carrier.

In classical probabilistic treatments, it is assumed that, under certain evidential conditions $\mathscr{E}$,[1] the value $\mathbf{Pr}(p|\mathscr{E})$ of the likelihood of a particular statement p is known. This view of evidence, while adequate to represent the informational conditions of most controlled experimental setups, fails, however, to adequately model the effects that acquiring similar information has on our state of knowledge when the state of the world can not be so readily manipulated.

In such circumstances, whenever the evidence $\mathscr{E}$ is observed, three possible informational outcomes may result from examination of further information that later turns out to improve our state of knowledge: either p is found to be true, $\neg p$ is found to be true (i.e., p is false), or such information is insufficient to determine the truth value of p. Use of modal logic concepts, which are the bases of the formal model of Ruspini [17], suggests the use of the notation $\mathbf{K}p$, $\mathbf{K}\neg p$, and $\mathbf{I}p$ to identify these outcomes. Since these alternatives are exclusive, it is clear that

$$\mathbf{Pr}(\mathbf{K}p) + \mathbf{Pr}(\mathbf{K}\neg p) + \mathbf{Pr}(\mathbf{I}p) = 1\,.$$

Furthermore, since the probability of $\mathbf{I}p$ may be positive, it will be true, in general, that

$$\mathbf{Pr}(\mathbf{K}p) + \mathbf{Pr}(\mathbf{K}\neg p) \leq 1\,.$$

This model, based on a combination of classical probability methods and the modal logic **S5** [8,12], essentially provides—through the logical notion of possible world— a meaning

[1]Throughout this paper, the symbol $\mathscr{E}$ is used to denote available evidence, i.e., a collection of propositions about the real world that are known to be true either as the result of direct observation or as the consequences of applicable background knowledge.

for the unary operator **K** as the representation of the state of knowledge of a statistician that is estimating the probability of truth of diverse propositions $\{p, q, \ldots\}$ under evidential conditions $\mathscr{E}$.

This statistician estimates those distributions by considering multiple samples of the state or behavior of a real-world system. Using, for each sample, additional information collected through further experimentation, the statistician may then establish or not the validity of a proposition p. If he is rather lucky, our statistician will find himself in the ideal situation where he can actually "know"[2] or "prove" that the real world is in a state s that is described to the best level of detail that is necessary to understand its behavior (i.e., a "possible world"). This is the state of knowledge usually attained, under perfect laboratory conditions, when experimental samples are fully analyzed and when the outcome of such analyses is classified in terms of a set of exhaustive and mutually exclusive alternatives.

Under less desirable epistemological circumstances, however, the statistician will only be able to prove that a less specific proposition p is true. In the extreme case where no further information exists, he will be forced to say that his knowledge is limited to that provided by the evidence $\mathscr{E}$, or that it is "vacuous."

All samples so analyzed, however, can be classified as to the "most specific knowledge" that could be determined in each case. The corresponding probability measure of the set $\mathbf{e}(p)$ of samples where the proposition p was the most specific knowledge (caled an *epistemic set* by Ruspini) corresponds, in Shafer's framework, to the value $m(p)$ of a mass function m, i.e.,

$$m(p) = \mathbf{Pr}(\mathbf{e}(p)).$$

Correspondingly, the probability that p was "known" to be true during statistical experimentation, corresponds to the value $\mathbf{Bel}(p)$ of Shafer's belief function, i.e.,

$$\text{Bel}\,(p) = \mathbf{Pr}(\mathbf{K}\,p).$$

The connection between the ability of our statistician to know that p was true and the belief and mass functions that he estimates through experimentation justifies both tthe expression *epistemic probability* introduced by Ruspini [17] to describe the underlying probabilities defined over a particular set of situations or scenarios $\mathbf{K}p$, (called the *epistemic universe*), and their description as being "probabilities of provability" or "probabilities of necessity" by Pearl [14], following a suggestion by Fagin and Halpern [6].

In short, all such interpretations are equivalent to the original model of Ruspini, where a rational agent was able to prove the truth of different propositions under different informational circumstances that were found to prevail, during his statistical experiment, with

---

[2] Note that, in the context of epistemic logics such as **S5**, the operator **K** behaves as a logical necessity operator. "Knowing" a proposition simply means that observations logically imply such proposition, or that it is *necessarily* true.

different frequencies of occurrence. Note, however, that while use of the terms "knowability," "provability," and "necessity" does much to provide adequate semantics to the calculus of evidence, its loose usage leads to unnecessary confusion. For example, in his recent criticism [15], Pearl takes some questionable semantic license with the term "necessity" mentioning, for example, the probability that a decision "will have to made out of compelling necessity." Such "pragmatic" necessity does not have anything to do, of course, with the "logical necessity" that underlies the Dempster-Shafer theory, i.e., the necessary truth of a proposition given available evidence.

Since the ability to prove a proposition q entails the ability to prove any proposition p that is implied by q, it should be clear that

$$\mathbf{Bel}(p) = \sum_{q \Rightarrow p} m(q)\,.$$

i.e., the fundamental equation relating the basic structures of the calculus of evidence. It is also true, as shown by Ruspini. that

$$\mathbf{Bel}(p) \leq \mathbf{Pr}(p) \leq 1 - \mathbf{Bel}(\neg p)\,.$$

providing bounds for the probability of p that may not be improved. This ability to manipulate probability intervals by means of the compact representation scheme of mass functions is the major reason for the appeal of the Dempster-Shafer methodology.

While the above discussion clarifies the nature of the statistician's knowledge modeled by belief and mass functions. doubts might still remain as to their utility to those that were not involved in their statistical estimation process. Such usage is, however, that made of any other probabilistic information. The analyst that observes $\mathscr{S}$ does not have the luxury that was available to the statistician estimating epistemic probabilities. i.e., the ability to collect additional information that permits a more detailed characterization of the state of the world, for the same reasons that the user of statistical tables is unable to utilize the raw data of the estimating statistician. Under such circumstances, the analyst is forced to rely on the probabilistic estimates provided by the statistician. which are believed on the basis of the assumed regularity of the repetitive behavior of the system: the epistemological cornerstone of probabilistic reasoning.

In other words, the "probability of provability" is the best information that is available to the analyst; an observation that not only disposes of questions about its role in probabilistic reasoning, but also of Pearl's worries about its use in lieu of the obviously more desirable "probability of truth" [15]:

> "why we should concern ourselves with the probability that the evidence implies A, rather than the probability that A is true, given the evidence?".

Clearly, we would prefer having the latter, but, unfortunately, we can only measure the former.

Our interpretation of the major evidential functions and structures also quickly disposes of erroneous arguments based on unintended interpretations of the intervals defined by belief functions. Each such interval represents ignorance of a single probability value for a proposition p under fixed evidential conditions $\mathcal{E}$. If critics choose, for example, to interpret such intervals as the possible values that conditional probabilities might attain when further evidence is collected, as suggested by Pearl [13], belief functions will not, indeed, behave according to such unintended semantics.

In closing this section, it is important to mention other alternative views of the structures of the calculus of evidence such as that recently proposed by Smets [19], which are based on a nonprobabilistic concept of belief. Although those models are interesting on the strength of their own virtues, we still emphasize that such interpretations are not required to reconcile the calculus of evidence with conventional probability theory.

In consideration of our ability to reconcile all structures and formulas of the calculus of evidence, including the Dempster's formula, with conventional probability structures, such as inner and outer probabilities, we do not feel strongly compelled to accept alternative epistemic interpretations. Our skepticism in this regard is further supported by the observation that, often, such epistemological alternatives are the result of misundertandings about the role of certain evidential formulas and processes (e.g., normalization). For the same reasons, we remain unconvinced about the need to assign several alternative interpretations to the structures of calculus of evidence or to its functions, as in the recently suggestion of Halpern and Fagin [7], which is echoed by Pearl [15].

3 On Decision Support

A criticism of a more fundamental nature of the calculus of evidence is often raised regarding the output of generalized interval-probability approaches. Since these methods often fail, due to basic knowledge deficiencies, to rank decision choices by the value of some measure that quantifies the desirability of each choice (e.g., expected utility), then it is said that they lack a decision-theoretic apparatus.

Although these arguments correctly point to the basic knowledge requirement that most decision problems entail--if a rational choice is to be made, then we must have a proper informational basis to do it— this obvious consideration is twisted, however, to argue for the necessity to estimate unknown probability and utility values when they are not available. We do not think that this *pragmatic necessity*, argument is either sound or compelling.

In our view, the calculus of evidence may be used in a straightforward fashion to produce intervals of possible utility-values. When such intervals overlap and cannot be ordered, this fact simply reflects a basic deficiency in our knowledge. We look down upon "pragmatic justifications" with the same concern that any experimental scientist must show about proposals to guess what he has not measured: the ability to make decisions in the absence of knowledge is, in our view, a handicap rather than an advantage of any method.

Far from lacking a decision-theoretic methodology, our approach provides an easily understandable quantification of the undesirable effects that poor information has on our decision-making ability; ordering decisions whenever it is rationally possible but advising us that such ranking is not possible if our knowledge is insufficient. In brief, our approach does not only supports decision-making but, through its built-in sensitivity-analysis features, helps us to determine what must be done to reach a happier epistemological state.[3]

4 On Dempster's Rule of Combination

The semantic model of the Dempster-Shafer theory also validates the so-called Dempster's rule of combination, which permits the combination of belief and mass functions corresponding to different evidential observations, made under certain conditions of independence. When such conditions are not valid, use of this formula leads, of course, to erroneous results, often, although incorrectly, considered to be an essential handicap of the evidential reasoning approach, rather than a consequence of its misapplication.

The Dempster formula is, currently, the principal evidence integration mechanism of the belief-function approach. It was derived in the context of a basic model of the effect of probabilistic evidence that correctly interprets such evidence as constraints on probability values rather than as the source of the actual values, which are typically undetermined. It may be described as an expression that, under certain conditions of independence, yields bounds for the conditional probability distribution $\mathbf{Pr}(\cdot|\mathcal{E}_1, \mathcal{E}_2)$ on the basis of similar bounds for the probability distributions $\mathbf{Pr}(\cdot|\mathcal{E}_1)$ and $\mathbf{Pr}(\cdot|\mathcal{E}_2)$.

To understand the conceptual bases for the Dempster's formula of combination and its consistence with conventional probability, we resort to a generalization of the logical model used before to derive the basic relations of the calculus of evidence. Instead of considering a single epistemic operator, corresponding to a single statistician or observer, we will consider two such rational agents, with their knowledge modeled by means of two operators $\mathbf{K}_1$ and $\mathbf{K}_2$. Each of these rational agents will be assumed to be ignorant of the knowledge possesed by the other, i.e., as if they were statisticians performing independent experiments under different evidential conditions $\mathcal{E}_1$ and $\mathcal{E}_2$. Their common knowledge, however, will be modeled by means of a nonindexed operator $\mathbf{K}$ corresponding to a third reliable agent that aggregates the statistical knowledge gathered by the other two.

Clearly, in a given applicable situation (i.e., the first agent observes $\mathcal{E}_1$ and the second agent observes $\mathcal{E}_2$), the integrating agent, who does not add any knowledge of his own, will be able to prove (or to "know" the truth of) a proposition p, if the other agents provide individual items of information that, when combined (i.e., conjoined) imply p, as expressed by the basic combination axiom:

[3]For an example of an approach that incorporates decision-maker preferences into the framework of the belief-function calculus, the reader is referred to a recent paper by Strat [21].

$\mathbf{K}p$ is true if and only if there exist sentences p_1 and p_2 such that $\mathbf{K}_1 p_1$ and $\mathbf{K}_2 p_2$ are true, and such that $p_1 \wedge p_2 \Rightarrow p$.

Using our three operators to generate all possible (i.e., logically consistent) states of knowledge that may be attained by each of the three agents while assessing the state of a real system, we may say that each of them has, as was the case before, a knowledge about the real world that may be represented by the "most specific"[4] propositions p_1, p_2 and p that each has been able to prove (with p being obviously more specific than either p_1 or p_2). In the terminology of Ruspini's semantic model, each of the agents is in an epistemic state, denoted by $\mathbf{e}(p)$, $\mathbf{e}_1(p_1)$ and $\mathbf{e}_2(p_2)$, respectively, each corresponding to the set of all conceivable states of the real world (i.e., possible worlds) having such knowledge characteristics.

The following important set-equation relating all of these types of epistemic sets as subsets of our enhanced epistemic universe, is the basis for the derivation of various evidential combination formulas

$$\mathbf{e}(p) = \bigcup_{p_1 \wedge p_2 = p} (\mathbf{e}_1(p_1) \cap \mathbf{e}_2(p_2)),$$

of which the Dempster combination formula

$$m(p) = \lambda \sum_{p_1 \wedge p_2 = p} m_1(p_1)\, m_2(p_2),$$

where

$$m(p) = \mathbf{Pr}(\mathbf{e}(p)|\mathscr{E}_1, \mathscr{E}_2), \quad m_1(p_1) = \mathbf{Pr}(\mathbf{e}_1(p_1)|\mathscr{E}_1), \quad m_2(p_2) = \mathbf{Pr}(\mathbf{e}_2(p_2)|\mathscr{E}_2),$$

and where λ is a multiplicative factor, is the best known and used.

Before reviewing the actual process leading to the derivation of the Dempster's formula, it is important to pause and reflect upon the nature of the above set-theoretic equation and its usefulness to derive evidence combination formulas.

We may first note that this equation has been derived as a relation between subsets of possible "epistemological states" that is valid regardless of any assumptions about probabilistic structures and their properties (e.g., independence). As such, it does not only provides the bases for the derivation of the Dempster formula but actually of a variety of formulas that bound possible probability values within and without the structures of the Dempster-Shafer theory.

Basically, this formula provides the basis to extend a probability function $\mathbf{Pr}$ that is known over subsets of the form $\mathbf{e}_1(p_1)$ and $\mathbf{e}_2(p_2)$ (i.e., over two *σ-algebras*), to the set of unions of sets of the form $\mathbf{e}_1(p_1) \cap \mathbf{e}_2(p_2)$ (i.e., another σ-algebra). If such extension can be made uniquely—as is the case for the Dempster formula—the resulting extension may be used to generate both the conditional probability $\mathbf{Pr}(\cdot|\mathscr{E}_1, \mathscr{E}_2)$ and its associated bounds Bel

[4]Note that such most specific knowledge always exists and is unique, but for logical equivalences, since the conjunction of all proved theorems is itself a theorem.

and Pl, which are fully compliant with the Shafer axioms. In other less fortunate cases (e.g., dependent evidence), such extension is not unique and the lower envelope of the possible extensions, which is not a probability, will lead to bounds that do not satisfy the axioms of the calculus of evidence.

A most important remark that must be made in this regard is that this equation is now being used to extend the evidential calculus approach by generalization of the notion of conditional probability by study of the probabilistic relations that define dependencies between the different types of epistemic sets (i.e., $\mathbf{e}(p)$, $\mathbf{e}_1(p_1)$ and $\mathbf{e}_2(p_2)$). Pearl [15], however, believes, apparently as the result of his examination of the role of *compatibility relations* in the calculus of evidence, that this approach is essentially limited in its expressive ability to set-theoretic relations between epistemic sets, which correspond to classical logical conditional statements (i.e., material implications).

In fact, it may be easily seen from our epistemic identity that whenever the conditional probabilities $\mathbf{Pr}(\mathbf{e}_2(p_2)|\mathbf{e}_1(p_1))$ and $\mathbf{Pr}(\mathbf{e}_1(p_1)|\mathbf{e}_2(p_2))$ are restricted to take the values 0 or 1,[5] then this identity may be used to map one body of evidence into another, i.e., by means of the compatibility relations that such probabilities define.

Since under these assumptions, however, there can be only one proposition p_2 for every proposition p_1 such that $\mathbf{Pr}(\mathbf{e}_2(p_2)|\mathbf{e}_1(p_1)) = 1$, and viceversa, then the compatibility relation that is so defined may be characterized by several implications of the form

$$\mathbf{e}_1(p_1) \Rightarrow \mathbf{e}_2(p_2)\,,$$

and of the form

$$\mathbf{e}_2(q_2) \Rightarrow \mathbf{e}_1(q_1)\,,$$

between knowledge states of one observer to knowledge states of the other which are useful to "transfer mass" between propositions. This correspondence must be contrasted with that following from the limited interpretation given by Pearl who, from knowledge of

$$\mathbf{e}_1(p_1) \Longrightarrow \mathbf{e}_2(p_2)\,,$$

concludes (by contraposition), correctly but narrowly, that

$$\neg\mathbf{e}_2(p_2) \Longrightarrow \neg\mathbf{e}_1(p_1)\,,$$

proceeding then to attach all material implication paradoxes (e.g., the "ravens paradox") to the calculus of evidence as if they were an essential methodological bane. If that were to be the case—clearly it is not— the same concerns should be raised about the use of conditionals in conventional probability calculus.

The second observation that may be made about the nature of evidence combination, in general, and the role of our basic set identity to generate combination formulas, in particular,

[5]It may be shown from the definition of epistemic sets that, under such conditions, knowledge of $\mathbf{Pr}(\mathbf{e}_2(p_2)|\mathbf{e}_1(p_1))$ suffices to derive $\mathbf{Pr}(\mathbf{e}_1(p_1)|\mathbf{e}_2(p_2))$.

is that while the functions to be combined are conditional probabilities over two different evidential sets $\mathscr{E}_1$ and $\mathscr{E}_2$, (i.e., the evidence observed by two agents), the desired integrated probability is a distribution over $\mathscr{E}_1 \cap \mathscr{E}_2$ (since we know that both observations are correct). Except for unusual cases, however. computation of $\mathbf{Pr}(\cdot|\mathscr{E}_1, \mathscr{E}_2)$ entails a "normalization" operation that is fully consistent with the calculus of probability. Most of the normalization "paradoxes" are the result of misunderstanding about what is being combined: two different conditional probabilities rather than two different lower and upper bounds of the same probability function.[6]

Focusing now on the rationale for Dempster's formula, we should notice first that the epistemic sets $\mathbf{e}_1(p_1)$ and $\mathbf{e}_2(p_2)$ are such that

$$\mathbf{e}_1(p_1) \subseteq \mathscr{E}_1, \qquad \mathbf{e}_2(p_2) \subseteq \mathscr{E}_2,$$

i.e., the possible knowledge states of each statistician include awareness of the truth of the evidence that is observed by each. Furthermore,

$$\mathscr{E}_1 = \bigcup_{p_1} \mathbf{e}_1(p_1), \qquad \mathscr{E}_2 = \bigcup_{p_2} \mathbf{e}_2(p_2),$$

where $p_1 \Rightarrow \mathscr{E}_1$ and $p_2 \Rightarrow \mathscr{E}_2$, i.e., each statistician knows something that implies that his evidential observation is true (otherwise he would not be "counting" that sample).[7]

Assume now that there exists a probability distribution $\mathbf{Pr}$ defined over the space of all possible epistemic states for our observing statisticians and our "integrating" agent. Each such epistemic state is a possible world that corresponds to a possible state of the world and to a possible state of knowledge for each agent that, in addition, is consistent with the laws of logic. We will assume now that, whenever $p_1 \Rightarrow \mathscr{E}_1$ and $p_2 \Rightarrow \mathscr{E}_2$, it is

$$\mathbf{Pr}(\mathbf{e}_1(p_1) \cap \mathbf{e}_2(p_2)) = \begin{cases} \mathbf{Pr}(\mathbf{e}_1(p_1))\,\mathbf{Pr}(\mathbf{e}_2(p_2)), & \text{if } p_1 \wedge p_2 \neq \emptyset, \\ 0, & \text{otherwise.} \end{cases}$$

This assumption simply states that, when $\mathscr{E}_1$ and $\mathscr{E}_2$ are both true the probability that a rational observer will be in a particular knowledge, or epistemic state does not provide any information about the probability of the epistemic state of the other agent (i.e., beyond ruling out logical impossibilities). In purely formal terms, we may say that knowledge of values of $\mathbf{Pr}$ over sets of the form $\mathbf{e}_1(p_1)$ does not provide any indication, beyond exclusion of logical impossibilities, of the values of $\mathbf{Pr}$ over sets of the form $\mathbf{e}_2(p_2)$ and viceversa. The epistemic states of our two agents may be said, therefore, to be uncorrelated in that knowledge of the state of one of our observers (by our integrating agent) does not provide any information about the state of the other, save for elimination of logical impossibilities.

[6]It is fair to say that much of the skepticism raised by the normalization used in Dempster's formula can be traced to the exposition given by Shafer [18], which suggests excessive reliance on unfounded heuristics.

[7]Recall that our observers, or rational agents, are statisticians estimating properties of certain statistical distributions by classifying each sample using their evidence and additional sample-dependent knowledge.

Noting now that

$$\mathbf{Pr}(\mathbf{e}_1(p_1)|\mathscr{E}_1) = \frac{\mathbf{Pr}(\mathbf{e}_1(p_1))}{\mathbf{Pr}(\mathscr{E}_1)}, \quad \mathbf{Pr}(\mathbf{e}_2(p_2)|\mathscr{E}_2) = \frac{\mathbf{Pr}(\mathbf{e}_2(p_2))}{\mathbf{Pr}(\mathscr{E}_2)},$$

$$\mathbf{Pr}(\mathbf{e}_1(p_1)\cap \mathbf{e}_2(p_2)|\mathscr{E}_1, \mathscr{E}_2) = \frac{\mathbf{Pr}(\mathbf{e}_1(p_1)\cap \mathbf{e}_2(p_2))}{\mathbf{Pr}(\mathscr{E}_1 \cap \mathscr{E}_2)},$$

then, whenever $p_1 \wedge p_2 \neq \emptyset$, it is

$$\mathbf{Pr}(\mathbf{e}_1(p_1)\cap \mathbf{e}_2(p_2)|\mathscr{E}_1, \mathscr{E}_2) = \lambda\, \mathbf{Pr}(\mathbf{e}_1(p_1)|\mathscr{E}_1)\, \mathbf{Pr}(\mathbf{e}_2(p_2)|\mathscr{E}_2) = \lambda\, m_1(p_1)\, m_2(p_2),$$

from which the Dempster's formula readily follows.

The normalization factor

$$\lambda = \frac{\mathbf{Pr}(\mathscr{E}_1)\,\mathbf{Pr}(\mathscr{E}_2)}{\mathbf{Pr}(\mathscr{E}_1 \cap \mathscr{E}_2)},$$

has been the object of considerable concern by both skeptics and proponents of the calculus of evidence. The above expression, however, provides the rationale for its usage while disposing of arguments about its alleged inconsistence with the probability calculus. In that expression, the denominator $\mathbf{Pr}(\mathscr{E}_1 \cap \mathscr{E}_2)$ appears as the consequence of the need to derive probability distribution estimates with respect to the intersection of the two observed evidences $\mathscr{E}_1$ and $\mathscr{E}_2$. The numerator of that expression simply reflects the need to combine conditional distributions over the same reference set (i.e., the epistemic universe) while our probabilistic knowledge is expressed over two of its subsets (i.e., $\mathscr{E}_1$ and $\mathscr{E}_2$).

The essence of the conditions that lend validity to the Dempster formula may be summarized by saying that its usefulness is confined to the limited, but rather important cases, where estimates of probabilistic likelihood have been formulated by two rational agents on the bases of independent observations, while ignoring the evidence available to each other.

If our integrating agent is thought of as being concerned with estimating the probabilities of certain events when both $\mathscr{E}_1$ and $\mathscr{E}_2$ are true, then we may say that, whenever the conditions validating the Dempster's formula hold, knowledge of the fact that a particular sample satisfies p_1, tells him nothing about the likelihood of p_2 (unless, of course, p_1 happens to be logically inconsistent with p_2). Furthermore, whenever our integrating agent is done with his job, he should find out that estimating this joint distribution (i.e., over $\mathscr{E}_1 \cap \mathscr{E}_2$) could have been accomplished in an easier fashion by estimating the marginal distributions over $\mathscr{E}_1$ and $\mathscr{E}_2$ and deriving the joint distribution by multiplication and normalization.

Other accounts supporting the validity of the Dempster's formula and its consistence with the probability calculus have been advanced by several authors. A particularly compelling justification has been recently given by Wilson [22].

5 On Paradoxes

Criticisms of the Dempster formula may be broadly characterized as being the consequence of basic misunderstandings about either its meaning or its validity.

In this section, we examine three alleged paradoxes of the theory showing that the purported inconsistencies are actually the results of conceptual misunderstandings or misrepresentations of the position of those who, while generally supporting the calculus of evidence, are concerned with its possible misapplication.

5.1 The Three Prisoner Problem

Turning our attention first to concerns about the validity of the Dempster's formula, we may note that, in general, such examples ignore its scope of applicability, producing counterintuitive results that are then used to dismiss the methodology as inadequate. Among those, the "three prisoner" problem discussed by Diaconis and Zabell [5] has been perhaps the more quoted and discussed.

This problem is one of a variety of examples, where the combination formula is used as a *conditioning formula* by assuming that one of the mass distributions being combined simply assigns all its mass to a proposition p in the frame of discernment. Combination of such simple support function with another mass function associated with a belief function $\text{Bel}(\cdot)$ leads to the conditioning formula

$$\text{Bel}(q|p) = \frac{\text{Bel}(q \vee \neg p) - \text{Bel}(\neg p)}{1 - \text{Bel}(\neg p)}.$$

In the particular case of the three prisoners problem, concerned with the guilt or innocence of a prisoner that has been chosen (by the Warden) as the guilty party by random draw among three candidates A_1, A_2 and A_3, our "logical space" or frame of discernment is simply the Boolean algebra induced by the three noncompatible propositions

"Prisoner A_i has been found guilty,"

where $i = 1, 2, 3$. Since only one of the three prisoners is chosen by the Warden, we clearly have

$$\mathbf{Pr}(p_i) = \tfrac{1}{3}, \quad i = 1, 2, 3$$

(Note that **Pr** is actually a conventional probability distribution).

Prisoner A_1 now asks the Jailer to name one of the innocent prisoners other than him arguing that such information would clearly be of little help to him as an indicator of his potential fate. As Pearl notes, if q stands for the proposition "The Jailer names A_2 as one of the innocent," then application of the conditioning rule leads to the result

$$\text{Bel}\,(p_1|q) = \text{Pl}\,(p_1|q) = \tfrac{1}{2}.$$

indicating that the conditional probability $\mathbf{Pr}(p_1|q)$ must be exactly $\frac{1}{2}$, instead of the "correct solution"

$$0 \leq \mathbf{Pr}(p_1|q) \leq \tfrac{1}{2}.$$

while also saying, agains the correct intuition of A_1 that his chances of guilt have been increased as the result of the irrelevant information provided by the Jailer. From such an observation, Pearl concludes that the formula is seriously flawed both because of the counterintuitive result that it produces and for its "collapsing" of a family of solutions into a single value.

Before proceeding to the discussion of Pearl's concerns we may note, in passing, that this problem has been well known as a source of paradoxes and incorrect solutions within the scope of the conventional probability calculus [2] quite independently of any issues of validity of its treatment using the Dempster-Shafer calculus. Curiously enough, the explanations given to describe the conceptual errors leading to incorrect classical treatments resemble to some extent that shedding light on the inapplicability of the Dempster's formula.

Returning now to the role of the Dempster's formula in this problem, we may first observe that, although, at first glance, the distributions representing the Jailer's and Warden's choices seem independent, it is actually impossible for the Jailer to tell to A_1 that A_2 is one of those to be spared if all he knew was that the Warden was choosing to be the guilty party by random draw (i.e., he needs to know exactly who is the one chosen for punishment). To use the terminology of Ruspini's model, the probability of A_2 being named as one of the innocent depends on the epistemic state of the Warden thus violating the independence assumptions of the Dempster's formula. If all possible combinations of truth values for the propositions p_i, $i = 1, 2, 3$, and q are tabulated, together with their probabilities, as done in Table 1, then it is clear that

$$\mathbf{Pr}(q|p_3) = 1\,. \quad \mathbf{Pr}(q) = \tfrac{1}{3}\,(1 + \alpha)\,.$$

where $0 \leq \alpha \leq 1$ represents the unknown probability that the Jailer will choose to name A_2 rather than A_3 as innocent if A_1 is actually the one chosen by the Warden as guilty.

Possible World	Warden's Choice	Jailer Identifies	Probability
W_1	A_1	A_2	$\frac{1}{3}\,\alpha$
W_2	A_1	A_3	$\frac{1}{3}\,(1-\alpha)$
W_3	A_2	A_3	$\frac{1}{3}$
W_4	A_3	A_2	$\frac{1}{3}$

Table 1: Possible Worlds in the Three Prisoners Problem

But then,

$$\mathbf{Pr}(q|p_3) \neq \mathbf{Pr}(q)\,.$$

violating the assumptions, discussed above, that validate the utilization of the Dempster's formula (i.e. $\mathbf{Pr}(\mathbf{e}_2(p_2)|\mathbf{e}_1(p_1)) \neq \mathbf{Pr}(\mathbf{e}_2(p_2))$). There is not, therefore, "total mistery," as

Pearl says, as to the incorrect results obtained using the Dempster's formula. Failing to be applicable, there should be little wonder that it leads to apparent paradox.

Although, as clearly shown by this discussion, the incorrect treatment of the three prisoner problem fails to invalidate the Dempster's rule of combination, we share the concern of Pearl and others about its wide misapplication, particularly when used undiscriminately to generate conditional distributions. In our research, we are endeavoring to extend the original theory to produce expressions to produce and utilize conditional belief information [16] that incorporates known dependencies between evidential bodies. These formulas are intended to provide better interval estimates that the typically uninformative bounds that are supplied by strict derivation of bounds in the absence of additional information by the expression

$$\mathrm{Bel}(q|p) = \frac{\mathrm{Bel}(p \wedge q)}{\mathrm{Bel}(p \wedge q) + \mathrm{Pl}(p \wedge \neg q)}$$

which is mentioned in Dempster's original paper [4] and that has been the object of recent concern by several authors [3,7].

In closing, we feel it is important to address other concerns of Pearl, going apparently beyond the three prisoners problem. about the counterintuitive nature of the "collapse" that usage of the Dempster formula often produces. which is manifested by production of a single conditional probability distribution when conditioning multiple members of a family $\mathcal{P}$ of probabilities over some specific subset q. Just as it is true that all members of the family of distributions

$$\mathcal{P} = \{\mathbf{Pr}_t : t \text{ in } [0,1]\}$$

defined in the set $X = \{a, b, c\}$ by the expression

$$\mathbf{Pr}_t(x) = \begin{cases} \frac{1}{2}t, & \text{if } x = a, \\ \frac{1}{2}(1-t), & \text{if } x = b, \\ \frac{1}{2}, & \text{if } x = c, \end{cases}$$

are such that $\mathbf{Pr}_t(\{a, b\}) = \frac{1}{2}$. despite their variability over other subsets. it is also true that an extensive family of distributions may collapse into a single conditional probability without violating any rational or probabilistic principles. Such "invariants" are, in fact, desirable as elements that simplify the analysis of an otherwise complex probabilistic problem. For these reasons, we do not feel that, if the Dempster's conditioning formula is applicable, its reduction of the variability of probability values should be a particular cause for concern as to its validity.

5.2 The Spoiled Sandwich

While discussing the suitability of the calculus of evidence either as a form of generalized probabilistic calculus, or as a new theory that intends to capture a novel notion of belief.

Pearl [15] again faults the approach for failing to satisfy the following rationality principle originally stated by Aleilunas [1]:

> "If two diametrically opposed assumptions yield two different degrees of belief in a proposition Q, then the unconditional degree of belief merited by Q should be somewhere between the two."

As natural such a principle might look at first, the following simple and clever example of Wilson [23] clearly shows that it is neither intuitive nor appealing pointing, however, to the pitfalls of creating or supporting one's favorite scheme on the strength of supposedly rational axioms.

Let $X = \{a,b,c,d\}$ with $A = \{a,b\}$ and $B = \{a,c\}$, so that $\bar{B} = \{b,d\}$. Consider the family of probability distributions in X

$$\mathcal{P} = \{\, \mathbf{Pr}_t : t \text{ in } [0,1] \,\},$$

indexed by a parameter t in $[0,1]$, and defined by

$$\begin{aligned} \mathbf{Pr}_t\,(\{a\}) &= \tfrac{1}{2}\,t, \\ \mathbf{Pr}_t\,(\{b\}) &= \tfrac{1}{2}\,(1-t), \\ \mathbf{Pr}_t\,(\{c\}) &= \tfrac{1}{4}, \\ \mathbf{Pr}_t\,(\{d\}) &= \tfrac{1}{4}, \end{aligned}$$

and let

$$\mathbf{Pr}_* = \inf_t \{\mathbf{Pr}_t\}\ .$$

Then, clearly

$$\mathbf{Pr}_t\,(A) = \tfrac{1}{2}t + \tfrac{1}{2}(1-t) = \tfrac{1}{2}\,,$$

and, therefore, it is $\mathbf{Pr}_*\,(A) = \frac{1}{2}$. The conditional probabilities $\mathbf{Pr}_t\,(A|B)$ and $\mathbf{Pr}_t\,(A|\bar{B})$ are given by the expressions

$$\begin{aligned} \mathbf{Pr}_t\,(A|B) &= \frac{\mathbf{Pr}_t\,(\{a\})}{\mathbf{Pr}_t\,(\{a,c\})} = \frac{\frac{1}{2}\,t}{\frac{1}{4}+\frac{1}{2}\,t}\,, \\ \mathbf{Pr}_t\,(A|\bar{B}) &= \frac{\mathbf{Pr}_t\,(\{b\})}{\mathbf{Pr}_t\,(\{b,d\})} = \frac{\frac{1}{2}\,(1-t)}{\frac{1}{4}+\frac{1}{2}\,(1-t)}\,, \end{aligned}$$

from which the lower bounds

$$\mathbf{Pr}_*\,(A|B) = \inf_t\, \mathbf{Pr}_t\,(A|B) = 0\,,$$
$$\mathbf{Pr}_*\,(A|\bar{B}) = \inf_t\, \mathbf{Pr}_t\,(A|\bar{B}) = 0\,,$$

are easily derived. It is clear, however, that

$$\tfrac{1}{2} = \mathbf{Pr}_*\,(A) > \mathbf{Pr}_*\,(A|B) = \mathbf{Pr}_*\,(A|\bar{B}) = 0\,,$$

showing that the the sandwich principle is violated even within the confines of conventional probability theory.

5.3 Other ways to spoil the sandwich

Although such simple examples should suffice to dispose of concerns about spoiled sandwiches, we feel that Pearl's discussion of the problem deserves a more detailed analysis, mainly because of its philosophical implications to rational thinking. This is particularly important as loose use of such terms as "assured winnings," "support," or "belief" in the absence of a sound formal interpretive framework may quickly mislead those engaged in the comparison of alternative methodologies.

In an example, called "the Peter, Paul, and Mary Sandwich problem," Pearl presents a betting situation where Mary prepares either a ham or a turkey sandwich promising to pay Paul $1000 should he guess correctly the type of sandwich that she has prepared. Not having a clue as to Mary's choice, Paul then flips a coin guessing "ham" if the coin turns up heads and guessing "turkey" if it comes up tails. Paul, as Pearl notes, behaves like an "incurable Bayesian," reckoning that

$$\begin{aligned}\mathbf{Pr}(\text{win}) &= \mathbf{Pr}(\text{win} \mid \text{turkey})\,\mathbf{Pr}(\text{turkey}) + \mathbf{Pr}(\text{win} \mid \text{ham})\,\mathbf{Pr}(\text{ham}) \\ &= \mathbf{Pr}(\text{tails} \mid \text{turkey})\,\alpha + \mathbf{Pr}(\text{heads} \mid \text{ham})\,(1-\alpha) = \tfrac{1}{2},\end{aligned}$$

regardless of the value α of the probability that Mary has actually prepared a turkey sandwich. Thus, in spite of not being "assured" a win, or having "supporting evidence," Paul can invoke the rationality (doubtful, as we already saw) of the sandwich principle and argue that Paul does not need to engage in unnecessary knowledge acquisition or experimentation [15]:

> "If every possible outcome of an experiment would lead you to choose the same action, then you ought to choose that action without running the experiment."

From such an observation, Pearl proceeds to fault the philosophical underpinnings of the belief-function approach eventually going as far as to suggest that, should Bayesian orthodoxy be unapplicable, the Dempster's formula—which, he freely admits, does not play any role in this example— be replaced by other formulas such as the well-known bounds recently rediscovered by Halpern and Fagin [7].

In the light of our previous example about the rather inconvenient ability of conventional probability families to spoil sandwiches, all of these pronouncements look increasingly suspicious: What, however, may we say that it is wrong? This question may be answered in two equivalent ways.

We may say first, keeping ourselves at the informal discussion level, that, often, the experiments may interact with probabilities in complex ways that, obviously, Pearl has not considered. Nothing in Pearl's formalism suggests, for example, that the sandwich has already been prepared and that it may not be artfully substituted by Mary to assure that Paul always loses thus invalidating his hopes of having at least a 50% chance of winning.

The second, more formal, rendering of this observation is again based on the semantic model of Ruspini. In this, and in other similar problems, we have several agents that deliberate about the state of the world on the basis of their knowledge and knowledge of the knowledge of others. If the unary operator $\mathbf{K}$ represents the state of knowledge of one of these agents, then, as observed before, our agent is usually in one of three possible epistemological states with respect to the validity of a proposition p: either he knows that p is true (denoted $\mathbf{K}p$), or he knows that p is false (denoted $\mathbf{K}\neg p$), or he may be ignorant of such truth (i.e., $\neg\mathbf{K}p \wedge \neg\mathbf{K}\neg p$, denoted $\mathbf{I}q$).

In standard accounts, assuming that knowledge of the truth of does not affect likehilood of truth of other propositions, we are simply concerned with a single form of conditional probability: that measuring the likelihood of p being true when q is true. In more complex epistemological situations, we may need to be concerned with such quantities as $\mathbf{Pr}(\mathbf{K}p \mid \mathbf{K}q)$, $\mathbf{Pr}(\mathbf{K}p \mid q)$, $\mathbf{Pr}(\mathbf{K}p \mid \mathbf{I}q)$, and the like. In other words, $\mathrm{Bel}(p \mid q)$ measures the support that knowledge of the truth of q provides to the truth of p, rather than the support provided by the truth of q to the truth of p.

In the Peter, Paul, and Mary sandwich problem, Pearl implicitly assumes that

$$\begin{aligned}
\mathbf{Pr}(\mathbf{K}_{\mathrm{MARY}}\mathrm{heads}) &= 0\,,\\
\mathbf{Pr}(\mathbf{K}_{\mathrm{MARY}}\mathrm{tails}) &= 0\,,\\
\mathbf{Pr}(\mathrm{turkey} \mid \mathbf{I}_{\mathrm{MARY}}\mathrm{heads}) &= \alpha\,,\\
\mathbf{Pr}(\mathrm{ham} \mid \mathbf{I}_{\mathrm{MARY}}\mathrm{heads}) &= 1-\alpha\,,
\end{aligned}$$

concluding correctly, by application of the total probability law, over the exhaustive and exclusive set of possibilities

$$\{\mathbf{K}_{\mathrm{MARY}}\mathrm{heads},\ \mathbf{K}_{\mathrm{MARY}}\mathrm{tails},\ \mathbf{I}_{\mathrm{MARY}}\mathrm{heads}\}\,,$$

that Paul has at least a 50% chance of winning.

This correct use of the total probability law does not mean that, by contrast, one should assume that the full extent of the conditional information provided by belief functions is limited to the conditional support functions

$$\mathrm{Bel}\,(p \mid q) = \mathbf{Pr}(p \mid \mathbf{K}q)\,, \quad \mathrm{Bel}\,(p \mid \neg q) = \mathbf{Pr}(p \mid \mathbf{K}\neg q)\,,$$

as Pearl evidently does. In short, not knowing p is not the same as knowing $\neg p$. The example of the Peter, Paul, and Mary sandwich shows that one needs to consider states of ignorance that, when properly accounted for, spoil even the best conceived principles of rationality.

To fully appreciate the complexity of the problem, suppose that we change Pearl's implicit assumptions bringing the previously absent Peter into the scene as a spy acting on behalf of Mary. In this new scenario, still consistent with Pearl's explicit statement of the problem, Peter, spying on Paul's coin flipping experiment, alerts Mary who, being rather artful and deft of hand, substitutes the sandwich so as to make sure that Paul always loses. In this case,

$$\mathbf{Pr}(\mathrm{ham} \mid \mathbf{K}_{\mathrm{MARY}}\mathrm{tails}) = 1\,, \quad \mathbf{Pr}(\mathrm{turkey} \mid \mathbf{K}_{\mathrm{MARY}}\mathrm{heads}) = 1\,,$$

and, most importantly,

$$\mathbf{Pr}\Big((\mathbf{K}_{\mathbf{MARY}}\text{heads}) \cup (\mathbf{K}_{\mathbf{MARY}}\text{tails})\Big) = 1\,,$$

i.e., Mary is never ignorant as to what Paul will bet.

The Peter, Paul, and Mary sandwich example does not, in our view, invalidate the applicability of the evidential approach but rather highlights the need to make necessary discriminations between propositional truth, knowledge of that truth, and the interplay between such conditions that are likely to be glossed over by cursory analyses based on conventional approaches.

5.4 The Disagreeing Experts

Another common misunderstanding regarding the role of the Dempster's combination formula is that provoked by an example of Zadeh [24] that, although originally formulated to illustrate problems with its misapplication, is often described as an indication of theoretical inadequacy.

Zadeh's example concerns two "experts" that assess, in a rather conflicting fashion, the likelihood of three, non-compatible, events A, B and C as shown in Table 2. Representation of each of the expert's assessment as a mass distribution followed by their combination with the Dempster's rule yields $\Pr(B) = 1$, indicating that the "true" event is B, an alternative considered to be rather unlikely by either of the assessors.

Observer	Pr(A)	Pr(B)	Pr(C)
1	0.99	0.01	0
2	0	0.01	0.99

Table 2: Experts Disagree on the State of the World

Although this example is often quoted as an example of the failure of the Dempster's rule, it is clear that each of the rows in Table 2 defines a conventional probability distribution, thus suggesting that the problem is likely to lie elsewhere. While one may be tempted to defend any method of evidence combination by saying that the evidence, however peculiar, indicates that Observer 1 is ruling out alternative C while Observer 2 is excluding alternative A, thus leaving only B as the sole possible answer, it is clear, upon further examination, that the rows of Table 2 cannot possibly be evaluations of the same probability distribution. If that were the case, then at least one of the experts must be wrong, since there can only be one correct probability distribution, contradicting the assumption that they are both reliable.

Clearly, if the example is to make any sense —under any type of probabilistic interpretation— each row must correspond to a different conditional probability where the conditions corresponds to different observations available to each expert. A simple example, suggested by a recent example used by Kyburg [9] to address other probabilistic reasoning issues, will help to clarify matters.

In this example we are being asked to reason, on the basis of available evidence, about the taste and edibility of certain berries that may be either small or large, red or blue, have good or bad taste, or be safe or poisonous to eat. We will assume that the berries in question are distributed according to the distribution shown in Table 3.

Berries			
Color	**Size**	**Taste/Edibility**	**Probability**
Red	Small	Good/Edible	99/199
Blue	Large	Bad/Edible	99/199
Red	Large	Poisonous	1/199

Table 3: The Berries Probability Distribution

If now a berry is picked up and found by an expert to be large, he will correctly conclude from such evidence that

$$\mathbf{Pr}(\text{Good}|\text{Large}) = 0\,, \quad \mathbf{Pr}(\text{Poisonous}|\text{Large}) = 0.01\,, \quad \mathbf{Pr}(\text{Bad Taste}|\text{Large}) = 0.99\,.$$

Another expert, noticing that the berry is red, will conclude, on the other hand, that

$$\mathbf{Pr}(\text{Good}|\text{Red}) = 0.99\,, \quad \mathbf{Pr}(\text{Poisonous}|\text{Red}) = 0.01\,, \quad \mathbf{Pr}(\text{Bad Taste}|\text{Large}) = 0\,.$$

Clearly the evidential implications of these two separate observations are identical to the situation summarized in Table 2. Examination of Table 3, however, reveals that

$$\mathbf{Pr}(\text{Poisonous}|\text{Red, Large}) = 1\,,$$

a correct solution that must be rationally be expected from any reasoning method that purports to be valid.

The solution to the puzzle of the disagreeing experts lies on recognizing that there is, in fact, no disparity of opinion among them. Each is providing quantitative measures of likelihood with respect to different reference classes. The Dempster formula, as observed by Zadeh, should never be applied to pool partial information about the same probability distribution. Furthermore, as shown by a sensitivity analysis of the results of its application to the berries example, its usage in situations where there is considerable disparity between reference classes (as suggested by the large normalization factor) should be discouraged on the basis of practical rather than conceptual considerations.

6 On Complexity and Generality

The potential complexity of the belief-function approach to represent and manipulate interval constraints on a family of probability distributions has been often mentioned as a handicap of the evidential reasoning methodology. In spite of such misgivings, two major empirical observations have indicated that the approach is applicable to a wide variety of practical problems.

First, our experience shows that, notwithstanding criticisms based on unrealistic worst-case scenarios, the approach is computationally efficient. In particular, we have found that representation of belief functions in terms of basic probabilistic assignments results in a storage and manipulation scheme that is both economical and easy to understand. In addition, we have sucessfully implemented tools, such as summarization and coarsening operators, which may be effectively utilized to limit representational complexity.

Second, our current functional operators have been chosen to guarantee that probabilistic information will always be capable of being represented within the scope of the approach, as more general constraints do not either enter into consideration or appear as the result of any of its functions.

The lack of generality of the belief-function approach to represent general lower-upper probability constraints is well known [10]. Our reliance on the methodology is primarily the result of practical considerations: while we would prefer to manipulate more general constraints on probability values, compelling computational efficiency arguments force us to limit the scope of the problems considered to those capable of being at least approximately solved by a belief-function treatment.

Being, in general, partial towards interpretations of evidential structures that are fully compatible with probability theory, our current research is being directed toward the development of more general, yet efficient, representation and manipulation methods.

Our current concerns with the manipulation of conditional and dependent evidence (i.e., the evidential counterpart of conditional probabilities) show, for example, that, for some important problems, the results of evidential combination fall outside the scope of its representational capabilities. In our experience, these methodological limitations are more worrisome than any of the supposedly paradoxical results arising from its misuse or its claimed lack of a decision-making apparatus.

Preliminary results [16] indicate, on the other hand, that the belief-function approach may be used to approximate the results of these evidential combination operations and that extended representation mechanisms [20] may yet be developed to treat more general evidential problems. This research also shows the basic errors inherent in criticisms that regard the belief-function approach as a fully developed methodology incapable of sustaining further enhancement and modification. Having been studied in depth for only fifteen years, its technological status is that of a young discipline being both capable of enhancement on its own and of combination with other approaches to produce more general tools for probabilistic reasoning. Far from proving that we have reached a technological plateau, our investigations indicate that much is yet to be gained from such a development and integration process.

7 Acknowledgment

The authors want to thank Aviv Bergman, Aaron Bobick, Piero Bonissone, Luis de Campos, Allan Chateauneuf, Mike Clarke, Didier Dubois, Tom Garvey, Joe Halpern, Andrew Hanson, Y.T. Hsia, Jean-Yves Jaffray, Henry Kyburg, María Teresa Lamata, Abe Mamdani, Serafín Moral, Henri Prade, Alessandro Saffiotti, Glen Shafer, Prakash Shenoy, Philippe Smets, Marcus Spies, Nic Wilson, Llorenç Valverde, Len Wesley, John Yen, and Lotfi Zadeh for useful exchanges and insights.

This work was supported in part by the Defense Advanced Research Project Agency under contract N00039-88-C-0248, in part by the Air Force Office of Scientific Research under Contract No. F49620-89-K-0001, and in part by the United States Army Research Office under Contract No. DAAL03-89-K-0156. The views, opinions and/or conclusions contained in this paper are those of the authors and should not be interpreted as representative of the official positions, decisions, or policies, either express or implied, of the supporting agencies or the United States Government.

References

[1] R.A. Aleilunas. A new normative theory of probabilistic logic. In *Proceedings of the 7th. Biennial Conference of the Canadian Society for Computational Studies of Intelligence (CSCSI-88)*, pp. 67–74. Edmonton, Alberta, 1988.

[2] Y. Bar-Hillel and R. Falk. Some teasers concerning conditional probabilities. *Cognition*, 11: 109–122, 1982.

[3] L.M. de Campos, M.T. Lamata, and S. Moral. The concept of conditional fuzzy measure. *International Journal of Intelligent Systems* (forthcoming).

[4] A.P. Dempster. Upper and lower probabilities induced by a multivalued mapping. *Annals of Statistics*, 38:325 339, 1967.

[5] P. Diaconis and S.L. Zabell. Some alternatives to Bayes's Rule. In B. Grofman and O. Guillermo, editors. *Information pooling and Group decision making*, pp. 25–38. JAI Press, Greenwich, Connecticut, 1986.

[6] R. Fagin and J.Y. Halpern. Uncertainty, belief and probability. Research Report RP. 619/(60901), IBM Almaden Research Center, Almaden, California, 1989.

[7] J. Halpern and R. Fagin. Updating beliefs versus combining beliefs. Unpublished report, 1989.

[8] G.E. Hughes and M.J. Creswell. *An Introduction to Modal Logic.* New York: Methuen, 1972.

[9] H.E. Kyburg. Jerry's berries. Unpublished report, 1985.

[10] H.E. Kyburg. Interval valued probabilities. In P. Smets, M.R.B. Clarke, C. Froidevaux and J. Campbell, editors, *Positions and Debates in Nonstandard Logics*, New York: Academic Press, forthcoming.

[11] J. Lowrance, T. Strat and T. Garvey. A Framework for Evidential-Reasoning Systems. *Proc. National Conference on Artificial Intelligence*, 1986, pp. 896–903.

[12] R. Moore. *Reasoning about Knowledge and Action.* Technical Note 191, Menlo Park, California: SRI International, 1980.

[13] J. Pearl. On probability intervals. *International Journal of Approximate Reasoning*, 2:211-216, 1988.

[14] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference.* San Mateo, California: Morgan Kaufmann, 1988.

[15] J. Pearl. Reasoning with Belief Functions: An Analysis of Compatibility. *International Journal of Approximate Reasoning*, 4: 363-390, 1990.

[16] E.H. Ruspini. Approximate Deduction in Single Evidential Bodies. *Proceedings 1986 AAAI Uncertainty Workshop*, Philadelphia, 1986.

[17] E.H. Ruspini. *The Logical Foundations of Evidential Reasoning.* Technical Note No. 408, Artificial Intelligence Center, SRI International, Menlo Park, California, 1987.

[18] G. Shafer. *A Mathematical Theory of Evidence.* Princeton University Press, 1976.

[19] P. Smets. Belief functions. In P. Smets, M.R.B. Clarke, C. Froidevaux and J. Campbell, editors, *Positions and Debates in Nonstandard Logics*, New York: Academic Press, forthcoming.

[20] M. Spies. Combination of evidence with conditional objects and its application to cognitive modeling. To appear in I. Goodman and H. Nguyen (editors). *Conditional Logic in Expert Systems.* North Holland, Amsterdam, 1990.

[21] T.M. Strat. Decision analysis using belief functions. *International Journal of Approximate Reasoning.* 4: 391-418. 1990.

[22] N. Wilson. *Justification, Computational Complexity, and Generalization of the Dempster Combination Formula.* Unpublished report, 1989.

[23] N. Wilson. Private communication, 1990.

[24] L.A. Zadeh. *On the validity of Dempster's rule of combination of evidence.* Memorandum No. UCB/ERL M79/24. ElectronicsResearch Laboratory, University of California. Berkeley, California, 1979.