APPROXIMATE REASONING: PAST, PRESENT, FUTURE

Technical Note No. 492

June 27, 1990

By: Enrique H. Ruspini, Sr. Computer Scientist
Artificial Intelligence Center
Computing and Engineering Sciences Division

Published in Information Sciences Volume 57-58, pp 297-317 (1991)

This work was supported in part by the Air Force Office of Scientific Research under Contract 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 1

2 The Development of Approximate Reasoning 2

3 Possible-World Models 4

4 The Semantics of Approximate Reasoning 6

4.1 Probabilistic Reasoning 7

4.2 Generalized Probabilistic Reasoning 11

4.3 Possibilistic Reasoning 13

5 Looking ahead 17

BIBLIOGRAPHY 19

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 foundations 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 $p \to q$, 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.

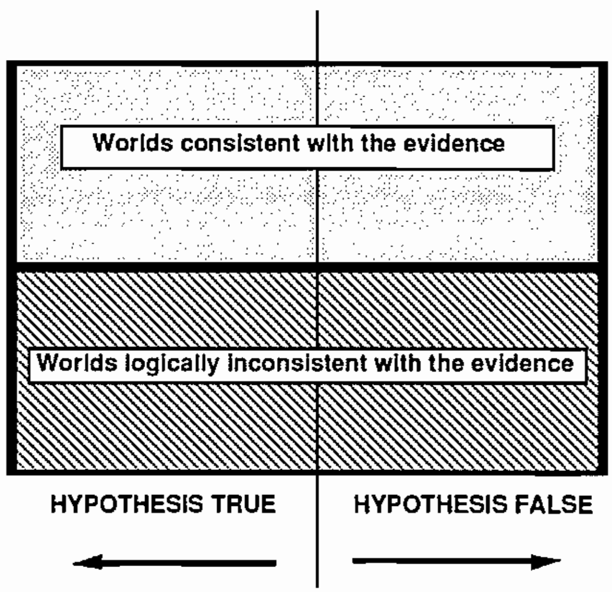

Figure 1: The approximate reasoning problem

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 e. 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 he 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*, *Lukasïewicz*, 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

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

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.

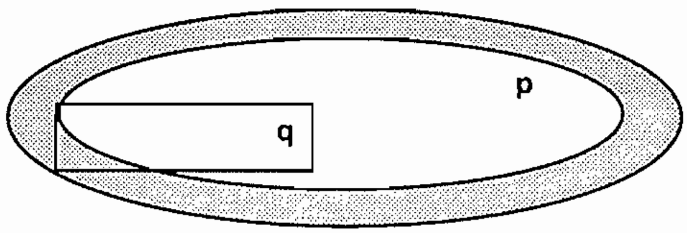

Figure 2: Degree of implication

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:

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

illustrated in Figure 3.

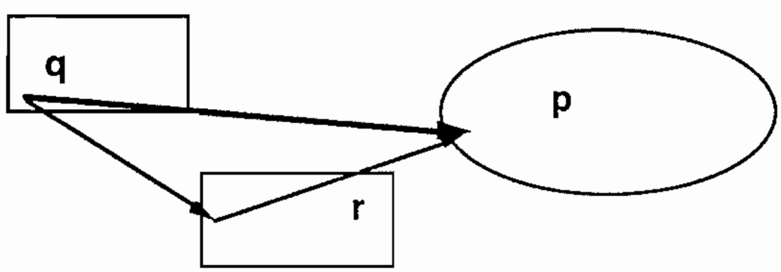

Figure 3: The generalized modus ponens.

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 Intelligence 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.