Responses to "An AI view of the treatment of uncertainty" by Alessandro Saffiotti

Dominic Clark (Editor), Jim Baldwin, Hamid Berenji, Paul Cohen, Didier Dubois, John Fox, John Lemmer, Henri Prade, David Spiegelhalter, Philippe Smets and Lotfi Zadeh

AI AND THE MANAGEMENT OF UNCERTAINTY

Dominic A. Clark

Imperial Cancer Research Fund Laboratories, 44 Lincoln's Inn Fields, London WC2A 3PX

1 Introduction

The management of uncertainty is a major part of many planning, decision making and problem solving tasks. Thus the development of robust techniques for managing uncertainty is a key factor in attempts to produce intelligent behaviour in machines.

The last few years have seen the development of a number of different formalisms for representing uncertainty and an ensuing vigorous argument about the management of uncertainty in AI systems (see Kanal and Lemmer, 1986). Many have come to suspect that the argument is now overly focused on technical aspects of the quantitative treatment of uncertainty, which may obscure deeper issues of greater long term practical significance. In particular, one of the fundamental demonstrations of AI seems to be ignored, i.e. *there is usually more than one representation for solving a problem*, and different representations have different uses and formal constraints on their use. It might be argued, therefore, that probability is an important representation of uncertainty, *but*

(1) it is only one representation;
(2) different representations may be needed on different occasions;
(3) different representations may be used simultaneously for different purposes within a single application, and
(4) machines should be able to reason about the representation(s) to be used in specific circumstances.

It is in this spirit that Saffiotti's (1988) wide ranging technical survey of prominent proposals takes an eclectic position which may help to settle an increasingly disruptive argument. Since the eclectic position has received relatively little attention, however, publication of Saffiotti's review was seen as an ideal opportunity to stimulate discussion of the proposal. The commentaries that follow are responses to this invitation.

This introduction provides an overview of Saffiotti's paper, four views across the commentaries and brief discussion of some emerging issues.

2 Commentary

2.1 Background

Saffiotti's paper is a review of the principal proposals for managing uncertainty in AI. It provides a general background and descriptions of Bayesian techniques, Dempster-Shafer belief functions (DS), Cohen's theory of Endorsements, other representation languages (multi-valued logics, fuzzy logics, implicit approaches to handling uncertainty), and the appropriateness of the various treatments. In the latter respect, his conclusion is. Namely, that different formalisms should be viewed as alternatives, and that the choice of an adequate technique for a particular application is a meta-level decision.

Discussion of the commentaries is structured around four themes: the eclectic position; the classification of uncertainty; extensions to, and refinements of Saffiotti's review; and more general comments.

2.2 The eclectic position

Many of the contributions explicitly supported an eclectic position, or acknowledged the limitations of a particular approach.

> "I believe that the identification of the type of uncertainty in a problem is an important factor in the selection of a proper method for its treatment." (Berenji)
>
> "... a consideration of the history of physics, philosophy, politics—and AI—will reveal that there is rarely one true representation of anything. Different representations are useful for solving different types of problems." (Fox)
>
> "... different representations may be needed on different occasions." (Lemmer)
>
> "I do believe that probability is not appropriate for representing all aspects of what is linguistically covered by the term uncertainty." (Spiegelhalter)
>
> "... the techniques in question ... have different domains of applicability, and ... a technique which is effective in relation to a particular dimension may be lacking in effectiveness along others." (Zadeh)

Other contributions agreed implicitly with the eclectic position by proposing situations in which the different management techniques may be more or less appropriate (Baldwin, Dubois and Prade).

2.3 Criteria of applicability

Several commentators proposed classifications of uncertainty relating to the applicability of the various management techniques.

Zadeh suggests that there are three principal non-independent facets of uncertainty relevant to AI which can occur in various guises and combinations. These are:

- Lexical imprecision or elasticity, as in "The price of oil is not *likely* to increase *substantially* in the *near* future".
- Granularity of information as in "Mary was born in March", where March plays the role of a crisp granule.
- Probability assignment and probability quantification, as in "the probability that Mary was born in March is 0.8".

Using this classification, he suggests that, in general, Bayesian methods are appropriate to those situations in which: (a) the premises involve probability assignments; (b) there is no lexical imprecision; and (c) there is no granularity of information. DS theory is appropriate where (a) and (b) are as above, but certain restricted types of granularity are allowed, and fuzzy logic is appropriate when lexical imprecision occurs. Baldwin asserts that uncertainty comes about as a result of some necessary form of incompleteness, specifically missing data or incomplete definitions. The former is amenable to representation by probability theory, the latter by fuzzy set theory.

Berenji suggests two aspects of uncertainty (fuzziness and ambiguity) as a basis for selecting an appropriate uncertainty method. Fuzziness stems from the lack of a clear boundary for the meaning of an expression (such as "Mary is *young*"). Ambiguity is the result of non-specificity of a term (as in "Mary is a *teenager*", where age might vary from 13 to 19), being a one-to-many relation.

He argues that multi-valued logics (and by implication, fuzzy logic) are more amenable to the representation of fuzziness while probability theory, possibility theory and DS theory are more amenable to the representation of ambiguity.

Dubois and Prade discuss the formal relationship between possibility theory, probability theory, Shafer's theory of evidence and Upper and Lower probability systems. They observe that probability theory can model pure randomness, but not total ignorance, while the converse is true for possibility theory.

Lemmer compares Bayesian Approaches with DS theory with respect to cost effectiveness of knowledge engineering and applicability. Regarding the former, he argues that:

> "it is well-known that eliciting probabilities is fraught with difficulties. Eliciting belief functions should be much more difficult still, not only because there are so many more, but also because their meaning is even less intuitive than the meaning of probabilities . . . while there are known serious hazards on the sea of Bayesian knowledge engineering, the sea of DS knowledge engineering is virtually uncharted."

With respect to applicability, Lemmer notes that for Bayes' rule to be applicable there must be adequate knowledge of a prior distribution. For DS, the bodies of evidence to be compared must be independent.

2.4 Extensions and refinements

A number of commentators noted recent developments to the methods reviewed by Saffiotti or gave details of systems omitted from his review.

Spiegelhalter pointed to recent developments in Bayesian theory whereby probabilistic relationships may be expressed graphically. This reduces some of the computational problems inherent in other Bayesian approaches.

Dubois and Prade highlight the distinction between fuzzy set theory and possibility theory. Essentially, although these can be represented within a common framework, fuzzy set theory is set-theoretic, while possibility theory is measure-theoretic.

Smets suggests an interpretation of DS theory that differs from the standard presentation, which he claims is a special case of upper and lower probabilities theory as proposed by Good.

Baldwin presents a further alternative to DS theory which he terms "support logic". This is a two valued calculus in which one value denotes the necessary support for a given hypothesis and the other denotes the necessary support against. It differs from DS theory in so far as an intersection rule is employed for combining support pairs rather than Dempster's rule of combination.

2.5 More general comments

Two commentators focus their comments on general aspects of uncertainty. Fox argues that: "the preoccupation with quantitative details of the representation, revision, and propagation of uncertainty has distracted us from understanding the limits of traditional mathematics theories of decision making and inhibited the development of new more versatile approaches." He draws an analogy between the contrasting qualitative and quantitative representations of space and time that have been employed in other areas of AI and the representation of uncertainty, and notes that: "If we can resolve our own differences we might shed some light outside our own parish [into other areas of AI]."

Paul Cohen's contribution takes the most radical approach to formal representation of uncertainty. His position can be summarized as follows:

- The analysis of uncertainty involves investigating how it arises, how it affects behaviour, and how we can minimise its impact.
- The literature in AI has been concerned almost exclusively with updating degrees of belief in interpretation tasks. Other generic tasks such as planning, design, robotics and process control have been largely neglected, thus most of the research published in the uncertainty literature is irrelevant to current AI research problems.

- The best AI work on uncertainty relegates the calculus to a detail, and solves the problem of uncertainty by control.
- An AI theory of uncertainty should be concerned with what one does in uncertain situations, and less concerned with how much one believes certain propositions. Thus, formal representations of uncertainty as degrees of belief are not sufficient (or even necessary) for problem solving under uncertainty.

3 Concluding remarks

Acceptance of eclectic position, discussed above, suggests that selection of an appropriate uncertainty management technique for a task might be viewed as a meta-level decision problem. With the corollary that a multi-task system might reason about the most appropriate technique for a particular task. Such a scenario is predicated, however, on the assumption that differential domains of applicability are identifiable for the various techniques, which, in turn, demands a comprehensive classification of types of uncertainty and of uncertainty management techniques.

While several of the commentaries have suggested partial classifications, it is clear that no consensus or complete scheme has emerged and a more wide ranging classification is required.

Bonissone (1987) has proposed 13 desiderata for a universally applicable uncertainty management scheme, which reflect putative deficiencies in one or more of the existing approaches (Table 1). These may also be viewed as dimensions or aspects of comparison between the various management procedures, and therefore constitute a basis for selecting among them.

Table 1 Desiderata for uncertainty management (Bonissone, 1987). (Emphasis added)

1. Combination rules *should not be based on global assumptions of evidence independence.*
2. The combination rules should *not assume the exhaustiveness and exclusiveness of the hypotheses.*
3. There should be an *explicit representation of the amount of evidence for and against each hypothesis.*
4. There should be an *explicit representation of the reasons for and against each hypothesis.*
5. The representation should allow the user to describe the uncertainty of any information at the available level of detail (i.e. allowing *heterogeneous information granularity*).
6. There should be an *explicit representation of consistency.*
7. There should be an *explicit representation of ignorance* to allow non-committal statements.
8. There should be a *clear distinction between a conflict in the information* (*violation of consistency*) *and ignorance* about the information.
9. There should be a *second order measure of uncertainty*, recording the uncertainty of the information as well as the uncertainty of the measure itself.
10. The representation must be, or appear to be, *natural to the user* to facilitate graceful interaction, natural to the expert to permit elicitation of consistent weights or reasons, and the semantics of procedures for propagating and summarising information must be clear.
11. The *syntax and semantics of the representation should be closed under the rules of combination.*
12. Making *pairwise comparisons of uncertainty should be feasible* as these are required for decision making.
13. The *traceability of the aggregation and propagation of uncertainty through the reasoning process* must be available to resolve conflicts or contradictions, to explain the support of conclusions, and to perform meta-reasoning for control.

However, classification of uncertainty based on these desiderata must be extended in at least two respects.

Firstly, there must be an explicit emphasis on cognitive aspects of uncertainty. Classification based on existing management techniques alone has the effect of obscuring those aspects of the uncertainty that are not well represented by existing formal models. Emphasis on the more cognitive aspects of uncertainty may suggest areas where sufficient management techniques have not yet been developed. A more cognitively based classification is provided by Mamdani et al. (1985, Table 2).

Table 2 Classification of uncertainty (Mamdani et al., 1985). (Emphasis added)

One's *belief* in a given proposition or a piece of knowledge.
The *likelihood* of a simple, compound or conditional event.
The *extent* of a proposition concerning a continuous variable.
The *imprecision* about any information.
The *exception* to a general rule.
The *mandate* for performing some action.
The *relevance* of one piece of information to another.

Secondly, the nature of the task in which the uncertain information is to be used must be considered for meta-level decision capabilities to be constructed. In systems that interact in a significant manner with users, for example, high levels of precision may be gratuitous, whereas precision may be appropriate as input to a statistical maximization procedures.

Overall, then, although Saffiotti's review and the ensuing commentaries suggest that the eclectic position is viable, it is clear that the demarcation of the various uncertainty management techniques, through a comprehensive classification of uncertainty and analysis of tasks involving uncertainty, remains a goal rather than an actuality.

4 References

Bonissone, P, 1987. "Reasoning plausible", in: *Encyclopedia of Artificial Intelligence*, Shapiro, S C Wiley (Eds.).

Kanal, L N and Lemmer, J F, (Eds.), 1986. *Uncertainty In Artificial Intelligence*, North-Holland.

Mamdani, A, Efstathiou, J and Pang, D, 1985. "Inference under uncertainty", in *Expert Systems 85*, Merry, M (Ed.), CUP.

Saffiotti, A, 1988. "An AI view of the teatment of Uncertainty", *Knowledge Engineering Review*, **2** (2).

UNCERTAINTY IN AI

J. F. Baldwin,

Engineering Mathematics Department, University of Bristol, England

Introduction

Saffiotti states that real world knowledge is accompanied by uncertainty. Constant conjunctions are much less important than frequent conjunctions in understanding the ways of the world, (Suppes, 1984). Causal relationships are probabilistic in nature and must be recognized when establishing valid causal laws, (Salmon, 1984). It would seem that certainty of knowledge in the sense of logical truth or perfect observation is not achievable.

This uncertainty comes about as a result of some necessary form of incompleteness. We should distinguish two forms of incompleteness:

(1) missing data
(2) incomplete definition of concepts

In the first case it is assumed that a logical conclusion could be reached if only some facts could be determined. Occlusion, for example, may prevent one concluding with certainty that some object in a photograph is some well-defined object. The truth that the man on trial is guilty of a given crime could be determined if certain evidence was available. This form of uncertainty is associated with probability theory.

If the crime is not well-defined this may not be possible even if the additional required evidence were available. This would correspond to the second case. Most concepts that we use, other than mathematically defined ones, are difficult to define in terms of necessary and sufficient conditions. What is meant by a humane society? When is a system stable? What do we mean by a safe nuclear plant? The world is continuous in nature but our understanding of it is in terms of discrete type concepts whose applicability cannot always be completely determined. As in the law courts, we often argue a case with reference to similar examples from previous experience. Fuzzy set theory is applicable to this type of uncertainty.

Human perception can be understood in terms of changes in neural activities with changing environment. These patterns of neural activities are approximately interpreted in terms of higher order concepts and conceptual relations. Uncertainty of the first type given above can be present in the neural responses while both types of uncertainty are relevant to the higher order mental interpretations. Both types of uncertainty are relevant to language and set theoretic semantics is inadequate for language interpretation, (Miller and Johnson-Laird, 1976).

Confirmation theories

Saffiotti discusses the various models of the axioms of probability theory as given by Kolmogarov. He does not accept the logic interpretation as a model since evidence which supports a given hypothesis to a certain degree may not provide any support against the hypothesis. This is relevant to the Shafer/Dempster theory of evidential reasoning. It is also the assumption made by Baldwin (1985, 1986, 1987) in his support logic programming. The use of a single number to provide the degree of confirmation for a given hypothesis in the light of certain evidence seems inadequate. This should be replaced by two numbers; one representing a necessary support for, $Sn(H \mid E)$ and one representing the necessary support against the hypothesis given the evidence, $Sn(\text{NOT } H \mid E)$. In Baldwin's theory of support logic, these form a support pair $[Sn(H \mid E), 1 - Sn(\text{NOT } H \mid E)]$. This support pair defines an interval which contains the point value probability $Pr(H \mid E)$. The evidence E is not sufficient to determine this point value precisely. The calculus used for combining the support pairs for compound statements must be consistent with the probability axioms for manipulating the $Pr(H \mid E)$ contained within the interval defined by the support pair. The theory of support logic is not equivalent to the Dempster/Shafer theory. An intersection rule for combining support pairs from different proof paths is used instead of the Dempster rule.

Heuristics used by experts are probabilistic in nature. Truth is not guaranteed when certain conditions hold. Difficulties of entailment when true propositions are replaced by highly probable propositions are well-known. Contraposition is valid for deductive entailments but does not hold for high probabilities. Deductive entailment is transitive but strong inductive support is not. More importantly the following valid argument of deductive logic does not carry over into the inductive case. If A entails B then A *and* C entails B. In fact if $Pr(B \mid A)$ is high, this provides no constraint on the value of $Pr(B \mid A, C)$ since

$$Pr(B \mid A, C) = \frac{Pr(C \mid A, B) . Pr(B \mid A)}{Pr(C \mid A)}$$

This, of course, is obvious from sample space considerations. All relevant criteria must be considered when giving supports to predicates as suggested by Hempel's maximum specificity conditions, (Hempel, 1968).

Constraint reasoning

Probabilistic reasoning can be viewed as constraint reasoning in which the various probabilistic statements given, provide evidence to constrain the probability of another statement to be contained within a certain interval.

If we know that:

$$\Pr(P > Q) = 2/3$$
$$\Pr(P) = 4/5$$

what can we conclude about Pr(Q)? A point value probability cannot be determined since the above two probabilistic statements give insufficient information for this. The statements constrain the interval which contains Pr(Q). Three possible cases must be analyzed, since the case corresponding to NOT(P → Q) AND (NOT P) is not possible because fo the inconsistency of the two statements.

Case 1	Case 2	Case 3
P → Q	NOT(P → Q)	P → Q
P	P	NOT P
Q	NOT P	world 1: Q
		world 2: NOT Q
Q: {1, 1}	Q: {0, 0}	Q: {0, 1}
x1	x2	x3

where xi is Pr(case i) and {a, b} means that NEC(Q) = a, POS(Q) = b, where a = 0 if NEC(Q) is false and 1 if it is true, and b = 0 if POS(Q) is false and 1 if it is true. NEC and POS are modal logic operators. Then since

$$x1 + x2 + x3 = 1$$
$$x1 + x2 = 4/5 \text{ since } \Pr(P) = \Pr(P \text{ AND } P \to Q) + \Pr(P \text{ AND NOT } (P \to Q))$$
$$x1 + x3 = 2/3 \text{ since } \Pr(P \to Q) = \Pr(P \to Q \text{ AND } P) + \Pr(P \to Q \text{ AND NOT } P)$$

so that

$$x1 = 7/15, x2 = 1/3, x3 = 1/5.$$

Hence Pr(NEC Q) = 7/15 and Pr(POS Q) = 7/15 + 1/5 = 2/3
From Pr(NEC Q) = <Pr(Q) = <Pr(POS Q), it follows that
Pr(Q) lies in the interval (7/15, 2/3).

This interval can be determined using linear programming formulations and this is discussed in Baldwin (ITRC, 1988).

Fuzzy sets

We associated incompleteness of definition with fuzzy set theory and possibility theory, (Zadeh, 1978). Fuzzy boundaries associated with concept definitions arise because of human or high-level interpretations of continuous phenomena in terms of discrete-like concepts. An interpretation of fuzzy sets can be given in terms of voting models (Baldwin, 1980, 1988).

A fuzzy subset f with respect to the set F is defined by means of a membership function Mf: F → [0, 1]. In other words an element, e, of the set F belongs to the fuzzy subset f with a degree of membership Mf(e). One possible interpretation is in terms of the voting behaviour of a population P, say, of persons, all of whom have their own understanding of the meaning of f. Let each person belonging to a population P vote on whether to accept or reject the membership of a given element e of set F as belonging to f. Each person must accept or reject, agree or not agree to the element's membership. Abstentions and partial agreements are not allowed. Mf(e) is equated to the proportion of persons who vote for accepting e as a member of f. Similarly (1 − Mf(e)) is the proportion of persons who reject e as belonging to f so that this is the membership level for e not belonging to f. Each individual will have threshold level such that if his doubt in an element e belonging to f increases above this level he will reject its membership, otherwise he will accept it. It is similar to a member of a jury having to say guilty or innocent for the person on trial. The evidence presented at the trial may not be conclusive but the person must still make a final judgement. If people are allowed to abstain we return to an analogue with support pairs.

In this voting model only the total vote is recorded and the actual voting of individual members is not known. In order to determine the voting of a compound statement given the overall voting scores for the elementary propositions of the statement, it is necessary to make assumptions about the voting behaviour of individuals of the population. Depending on these assumptions various calculuses can be derived. For conjunction we can define a mapping T

$$T:[0, 1]*[0, 1] \to [0, 1]$$

which satisfies the axioms

(1) $T(a, 1) = a$
(2) $T(a, b) = T(b, a)$
(3) $T(a, b) >= T(c, d)$ if $a >= c$ and $b >= d$
(4) $T(a, T(b, c)) = T(T(a, b), c)$

T is called a T-norm and generalizes the *and* corresponding to conjunction.

Examples of instances of T are

(1) $T(a, b) = \min\{a, b\}$
(2) $T(a, b) = a.b$
(3) $T(a, b) = \max\{a + b - 1, 0\}$

A dual norm, called the T-conorm, S, exists which generalizes disjunction. For any T-norm T there exists a dual norm S such that

$$S(a, b) = 1 - T((1 - a), (1 - b))$$

The mapping

$$S:[0, 1]*[0, 1] \to [0, 1]$$

satisfies the same axioms as for T except that (1) is replaced by (1′) where

$$(1') \quad S(a, 0) = a$$

Examples of instances of S conorms corresponding to the T norms (1), (2) and (3) above are respectively

(1) $S(a, b) = \max\{a, b\}$
(2) $S(a, b) = a + b - a.b$
(3) $S(a, b) = \min\{a + b, 1\}$

This treatment of interpreting fuzzy sets can be extended to provide a form of semantic unification which generalizes the unification of logic programming which is purely syntactic. A probabilistic interpretation given to X is f1 given that X is f2 where f1, f2 are fuzzy sets on F, and X is a relevant first order logic construct. Pr(X is f1 | X is f2) is interpreted in this voting model as Pr(P accepts X is f1 | P is told X is f2).

Semantic unification is part of the support logic programming theory as provided in the FRIL language (Baldwin et al., 1987).

References

Baldwin, J F, 1985. *Support Logic Programming*, Proc. of the NATO Advanced Study Institute on Fuzzy Sets Theory and Applications, Louvain-La-Neuve, Belgium.

Baldwin, J F, 1986. 'Support logic programming", *Fuzzy Sets Theory and Applications*, Jones A and Zimmermann H J (Eds.), pp. 133–170, D Reidel Pub. Co.

Baldwin, J F, 1987. "Evidential support logic programming", *Fuzzy Sets and Systems*, pp. 1–26.

Baldwin, J F, Martin, T P and Pilsworth, B W, 1987. Fril Manual, Equipu AIR Ltd, 184 Hotwell Rd, Bristol.
Baldwin, J F, 1988. "Computational models of uncertainty reasoning in expert systems", to appear in: *Journal of Computers and Mathematics with Applications.*
Baldwin, J F, 1988. *FRIL Support Logic Programming*, ITRC report, University of Bristol.
Hempel, C G, 1968. "Maximum specificity and lawlikeness in probabilistic explanation", *Philosophy of Science*, **35** 116–133.
Miller, G A, Johnson-Laird, P N, 1976. *Language and perception*, Harvard Univ. Press.
Salmon, W C, 1984. *Scientific Explanation and the Causal Structure of the World*, Princeton Univ. Press.
Suppes, P, 1984. *Probabilistic Metaphysics*, Blackwell.
Zadeh, L A, 1978. "Fuzzy sets as a basis for a theory of possibility", *Fuzzy Sets and Systems* **1** 3–28.

RESPONSE TO SAFFIOTTI'S "AN AI VIEW OF THE TREATMENT OF UNCERTAINTY"

Hamid R. Berenji
NASA Ames Research Center/Heer Associates Inc., Mail stop: 244-17, Moffett Field, CA 94040

Saffiotti has advocated an *open mind* policy toward choosing a proper scheme of uncertainty management based on the context in which the problem solving is taking place. He has adequately called these *parallel uncertainty inference* techniques which have many similarities to what some of us have been advocating for some time (see Berenji, 1987). In Thompson, 1985; Fox, 1986; Bonnissone and Decker, 1986; Berenji, 1987, we have called for *an integrated approach toward reasoning with uncertainty* where, depending on the type or types of uncertainty in the problem, different uncertainty management scheme(s) could be selected.

Since Saffiotti's paper covers a wide range of issues, and much of the material is quite old with a long history of heated debate, I will limit my comments to the points that I disagree with most. I also comment on the points that would complement Saffiotti's interesting paper. In addition, I will provide some additional signposts to literature which are of critical importance but unfortunately not mentioned in Saffiotti's work.

1 Multiple faces of uncertainty

In discussing the *thousand and one faces of uncertainty*, Saffiotti refers only to the various interpretations of probability (e.g. frequentist, subjectivist, etc.). Although he implicitly refers to other types of uncertainties in the background by reviewing their dedicated techniques (e.g. fuzzy logic), he does not make this distinction explicit. In this respect Fox (1986) has provided a more complete discussion. Although it is realized that uncertainty in general is due to the lack of knowledge about the precise meaning of a term, a further break down is necessary for identifying its source and its nature. I believe that the identification of the type of uncertainty in a problem is an important factor in the selection of a proper method for its treatment. Two types of uncertainty immediately come to mind. The first type is due to *fuzziness* and the second is due to *ambiguity*.[1]

Fuzziness stems from the lack of clear boundaries for the meaning of an expression (e.g. Mary is *young* or the pressure on a valve is *moderately high*). This concept has more similarities to the concepts used in multi-valued logics where terms can take truth values in a range (e.g. between 0 and 1). In real life, the truth values of most concepts do not change abruptly. A good example of this is the degree of *injury* that insurance companies use to calculate their payments. Another example is a carton of milk in a refrigerator which moves, over time, from being good (not spoiled) to being spoiled.[2] Although it might be possible to represent some of these vague concepts in terms of

[1] Note that *randomness* falls better under this type of uncertainty.
[2] Also discussed in (Cheeseman, 1988; Yager, 1988).

subjective probabilistic expressions (e.g. by designing a test which seeks a consensus of *n* people with regard to whether they classify a certain height as average or tall) the resulting representations are often unnatural and unnecessarily complex. I believe that most of the human processing of uncertain situations (e.g. common sense reasoning) uses elements with gradual truth values and vague terms; in a sense analog processing. Developing further techniques for addressing these issues is of crucial importance and should be included in the foreground of every general treatment of uncertainty in AI.

Ambiguity is mainly the result of *non-specificity* in expressing the meaning of a term (e.g. *Mary is a teenager*, where assuming integer values for age, her age might be any number between 13 and 19 inclusive). In general, ambiguity is associated with one-to-many relations. A typical example in this case is when we have a crisp set of hypotheses and we seek a measure to rank these elements based on some criteria (e.g. the most likely disease or the most probable location of an object, etc.). Probability theory, Dempster-Shafer theory, and Possibility theory deal with this type of uncertainty. In this respect, Klir (1988) provides a good discussion of the types of uncertainty and the different measures that can be used to identify the level of uncertainty at each stage of the problem solving process.

2 Computational vs epistemological adequacy

A controversial topic in reasoning with uncertainty is the amount of input data that each approach requires and the availability of that data. This has been referred to as *computational adequacy criticism* by Saffiotti. Arguments with regard to the unavailability of prior probabilities required for many Bayesian analyses fall under this criticism. On the other hand, *epistemological adequacy criticism* refers mainly to the form of representation used for modelling uncertainty and its subsequent treatment. In this respect, the use of one-valued measures (e.g. subjective probabilities), two-valued measures (e.g. necessity and possibility, belief and plausibility), or fuzzy-valued measures (e.g. linguistic terms) have been largely debated in recent AI literature. I believe that these two (i.e. computational vs epistemological adequacy) are inseparable and should always be jointly evaluated while selecting an uncertainty management scheme.

It is often easy to fall into philosophical discussions about whether probability could, or could not do the things that alternative approaches can do (Cheeseman, 1985; Lindley, 1988; Zadah, 1986; Shafer, 1987; Dubois and Prade, 1985). However, as mentioned by Saffiotti (and I agree completely) this often shifts the criticism from the epistemological adequacy to the computational adequacy. For example, the solution provided by Bayesians often requires further assessment of probabilities (e.g. Lindley's *extending the conversion method* or Cheeseman's extensive use of new conditional probabilities 1988). The point here is that one should always try to select an adequate approach for the problem based on the epistemological and computational analyses. A good practical example is to note the progress that has been made in recent years in the research on application of *fuzzy control*. By modelling the vague terms that human operators use in controlling the operation of complex non-linear systems, it has become possible to control many processes on a real-time basis where real-time ranges from fractions of a second in control of an aircraft (Larkin, 1985) to minutes in control of a cement kiln (Ostergaard, 1977).

On a slightly different but related note, recent developments in the theory of belief functions (Shafer, Shenoy, and Mellouli, 1987; Shenoy and Shafer, 1986; Kong, 1986) have resulted in efficient algorithms which reduce significantly the applicability of the *computational adequacy criticism* to this theory.

3 Additional signposts to literature

An important contribution to the Certainty Factor approach is the work of David Heckerman, 1985 in providing a probabilistic interpretation for the method. In this work, he discusses the importance of distinguishing between the *measure of absolute belief* in a hypothesis and the *measure of change in belief* about a hypothesis.[3]

Another very important and recent contribution to this field is Ruspini's work in providing a logical foundation for the belief functions theory (1986). In this work, he applies the method proposed by Carnap (1962) to structures based on epistemic logics. Epistemic logics are modal logics which deal with issues related to the state of knowledge of rational agents about the real world.

Bonnissone in a series of papers (1987) uses Triangular Norms and Triangular Co-norms for generalization of the combination operators in rule-based systems and applies them in a layered architecture for uncertainty management.

Finally, a recent issue of the *Computational Intelligence* includes a very up to date and interesting debate on issues dealing with uncertainty and logic, Bayesianism, etc. (Cheeseman, 1988).

In summary, I share Saffiotti's view with regard to an open mind policy toward the selection of a proper approach for uncertainty management in AI. However, I believe that the identification of the source and the nature of the uncertainty in the problem deserve more attention than that discussed by Saffiotti in his paper.

References

Baldwin, J F, 1986. "Automated fuzzy and probabilistic inference", *Fuzzy Sets and Systems* **18**, 219–235.

Barnett, J, 1981. "Computational methods for a mathematical theory of evidence", *Proceedings of IJCAI 81* pp. 868–875.

Berenji, H R, 1987. An integrated approach to reasoning with uncertainty in artificial intelligence", *Proceedings of the second NASA Ames AI Forum*, Palo Alto.

Berenji, H R, 1988. "Treatment of Uncertainty in Artificial Intelligence", in: *Machine Intelligence for aerospace robotics*, (eds., Ewald Heer and Henry Lum), AIAA.

Bonissone, P P and Decker K S, 1986. "Selecting uncertainty calculi and granularity: an experiment in trading-off precision and complexity", *Uncertainty in AI*, Kanal and Lemmer (Eds.), pp. 217–247, North-Holland.

Bonissone, P, 1987. Summarizing and propagating uncertain information with triangular norms", *International Journal of Approximate Reasoning*, **1**(1) pp. 71–101, North Holland Publishing Company.

Bonissone, P P, Gans, S S and Decker, K S, 1987. RUM: a layered architecture for reasoning with uncertainty", *IJCAI-87*, Milan, Italy.

Carnap, R, 1962. *Logical Foundations of Probability*, University of Chicago Press, Chicago, Illinois.

Cheesemen, P, 1988. "An inquiry into computer understanding", *the Computational Intelligence* **4**, pp. 58–66.

Cheeseman, P, 1988. In defense of an inquiry to computer understanding", *the Computational Intelligence* **4**, pp. 129–142.

Cheeseman, P, 1986. "Probabilistic vs Fuzzy reasoning", in: *Uncertainty in AI*, Kanal and Lemmer (Eds.).

Cheeseman, P, 1985. "In defense of probability", *Proceedings of IJCAI 85*, pp. 1002–1009.

Dubois, D and Prade, H, 1985. "Combination and propagation of certainty with belief functions – a re-examination", *proceedings of IJCAI-85*, Los Angeles, CA.

Fox, J, 1986. "Three arguments for extending the framework of probability", in: L N Kanal and J F Lemmer (Eds.), *Uncertainty in Artificial Intelligence*, Amsterdam: North Holland.

Gordon, J and Shortliffe, E, 1983. "The Dempster-Shafer theory of evidence", in: *Rule-Based Systems*, Buchanan, B and Shortliffe, E (Eds.), Menlo Park: Addison-Wesley.

Heckerman, D, 1985. "A probabilistic interpretation for MYCIN's certainty factors", in: *Proceedings of AAAI Workshop on Uncertainty and Probability in AI*, Los Angeles, CA, pp. 9–20.

Kong, A, 1986. *Multivariate Belief Functions and Graphical Models*, Research report S-107, Harvard University.

Klir, G J and Folger, T A, 1988. *Fuzzy Sets, Uncertainty, and Information*, Prentice Hall.

Larkin, L I, 1985. "A fuzzy logic controller for aircraft flight control", *Industrial Applications of Fuzzy Control*, Sugeno (Ed.).

Lindey, D V, 1971. "Making decisions", John Wiley, London.

Lindley, D V, 1987. The probability approach to the treatment of uncertainty in artificial intelligence and expert systems", *Statistical Science*.

[3] These two concepts were used interchangeably in the original formulation of this method by Shortliffe (1976).

Ostergaard, J J, 1977. "Fuzzy logic control of a heat exchanger process", in: *Fuzzy Automata and Decision Processes*, North-Holland.
Prade, Henry, 1985. "A computational approach to approximate and plausible reasoning with applications to expert systems", *IEEE Transactions on Pattern Analysis and Machine Intelligence* **PAM-7**(3).
Pearl, Judea, 1986. "Fusion, propagation, and structuring in belief networks", *Artificial Intelligence* **29** pp. 241–288.
Ruspini, E, 1986. *Logical Foundation of evidential reasoning*, SRI Tech. note 408.
Shafer G, 1987. "Probability judgment in artificial intelligence expert systems", *Statistical Science*.
Shafer, G, Shenoy, P P and Mellouli, K, 1987. "Propagating belief functions in qualitative markov trees", *International Journal of Approximate Reasoning* **1** pp. 349–400.
Shenoy, P P and Shafter, G, 1986. "Propagating belief functions with local computations", *IEEE Expert*, pp. 43–52.
Shortliffe, E, H, 1976. "Computer based medical consultations: MYCIN", *American Elsevier*, New York.
Thompson, T R, 1985. "Parallel formulation of evidential reasoning theories", *Proceedings of IJCAI-85*, Los Angeles, CA.
Yager, R, 1988. "Responses to Peter Cheeseman's An Inquiry into Computer Understanding", *Computational Intelligence*, pp. 125.
Zadeh, L A, 1986. "Is probability theory sufficient for dealing with uncertainty in AI: a negative view", in: *Uncertainty in AI*, Kanal and Lemmer (Eds.), North Holland.

COMMENTS ON SAFFIOTTI'S "AN AI VIEW OF THE TREATMENT OF UNCERTAINTY"

Paul R. Cohen

Experimental Knowledge Systems Laboratory, Department of Computer and Information Science, University of Massachusetts, Amherst, Mass 01004

Although I agree with many of Saffiotti's statements, I think his paper is conservative. He starts well when he says that we should first understand "what uncertainty really is", but, characteristically, he doesn't follow through; he doesn't look at how uncertainty arises, how it affects behaviour, and what we do to minimize its impact. Indeed, within the first page of his paper, Saffiotti slides into a discussion of *probability*, apparently without explaining the relationship between probability, which is a formalism, and uncertainty, which is a state of mind with causes, consequences and associated problem-solving strategies. If he had thought about "what uncertainty really is", he would not have concluded that uncertainty is probability, or he would at least have justified the association. Instead, he accepts Thompson's framework, going so far as to declare: "The updating mechanism is pivotal to any theory". Well, it's pivotal to any theory *of updating*, but it plays a much smaller role in AI systems. As I will show shortly, Saffiotti virtually ignores the role of uncertainty in AI systems, and the AI solutions that have been invented to handle it. Consequently, his paper is not, as advertised, "An AI view of the treatment of uncertainty".

The reason this matters is that Saffiotti is typical. The literature on uncertainty in AI is concerned almost exclusively with a single aspect of a single task: updating degrees of belief in interpretation tasks. Expert systems for medical diagnosis serve as the canonical AI systems. Other aspects of interpretation tasks, such as deciding which evidence to acquire and deciding how to treat a patient given a diagnosis, have been largely ignored. Other generic tasks and domains, such as planning, design, robotics, and process control, have also been neglected. But AI is increasingly concerned with decision making in such tasks and domains, and is much less concerned with judgment in simple diagnosis. Thus, most of the research published in the uncertainty literature is irrelevant to current AI research problems, even though these problems involve considerable uncertainty. For example, of approximately 50 papers published in the *Proceedings of the Third Annual Workshop on Uncertainty in Artificial Intelligence*, roughly one half mentioned no application whatsoever but were clearly influenced by diagnostic expert systems, two described vision systems, and all but four of the rest described diagnostic applications, or algorithms for learning or knowledge acquisition for diagnostic applications (Tong and Appelbaum, 1987; Hanks, 1987; Cohen, 1987; Wellman, 1987).

So what is the alternative? How else should we study uncertainty? Isn't it, as Saffiotti suggests, just a matter of selecting the right calculus? Not at all. From my perspective, the best AI work on reasoning under uncertainty relegates the calculus to a detail. I am thinking of Hearsay-II (Erman and Lesser, 1975; Erman et al., 1980), which solves interpretation problems under enormous uncertainty by carefully controlling its inferences (see also (Hayes-Roth, 1985) or, Clancey's NEOMYCIN (1984), which plans medical diagnoses; or our own MUM system that similarly plans a sequence of medical tests to ensure the most diagnostic workup at the least cost (Cohen et al., 1987). The important characteristic of these and other programs (see Cohen, 1987) for a survey is that they solve problems under considerable uncertainty by *control*, that is, by exploiting some aspects of the problem to determine the best order of problem solving actions. For example, Hearsay-II's opportunistic control strategy exploited redundancy in the speech signal. More recently, we have come to realize that *planning* tasks also require problem solvers to carefully consider the order of their actions in uncertain situations. Whereas early AI planners assumed complete, accurate knowledge about the state of the world and the effects of actions, these assumptions are clearly unrealistic and must be discarded in, say, robot planning, job shop scheduling, process control, and similar tasks. In fact, uncertainty is the principle problem in planning today (see (Cohen and Day, 1988) for a discussion of how uncertainty affects planning, and styles of planning to cope with uncertainty).

My own research in control and planning has emerged from earlier work on endorsements, cited in Saffiotti's paper: "A declared goal of Cohen's approach is to use knowledge about uncertainty to affect system behaviour (for instance, to decide what to do to reduce uncertainty"). Saffiotti correctly points out that we had little success combining endorsements, partly because the function of endorsements is less to represent degrees of belief than to carry information the system needs to act appropriately. In a series of papers, I evolved the position that, from an AI perspective, a theory of uncertainty should be a theory of *problem solving actions under uncertainty* – it should be less concerned with how much one believes propositions, and more concerned with what one does in uncertain situations (Cohen, 1987; Cohen et al., 1987; Cohen, Greenberg and Delisio, 1987; Cohen, 1987; Howe and Cohen; 1986; Cohen and Day, 1988). Control and planning provide two ways to look at action under uncertainty, but it is easy to envision other approaches, such as hedging bets, worst-case analysis, simulation, and so on. Currently, we study these issues in several tasks, including planning medical diagnoses, dynamic real-time planning to control forest fires, and planning competitive strategy in markets.

In sum, the goal of AI is to have computers solve problems, so the goal of AI and uncertainty research is to have computers solve problems that afford significant uncertainty. Saffiotti's paper says a lot about *representations* of uncertainty, but it is well-known in AI that no representation is epistemologically neutral – it always serves some purposes better than others. The representations Saffiotti discussed are, with few exceptions, appropriate for degrees of belief. They are not sufficient (or even necessary) for problem solving under uncertainty. Unless researchers in the uncertainty community take problem solving seriously, their work will continue to be largely irrelevant to AI.

References

Clancey, William J, 1984. "Acquiring, representing, and evaluating a competence model of diagnosis", KSL Memo 84-2, Stanford University, February 1984. To appear in *Contributions to the Nature of Expertise*, Chi, Glaser and Farr (Eds.), in preparation.

Cohen, Paul R, 1987. "The control of reasoning under uncertainty: a discussion of some programs", *The Knowledge Engineering Review* **2**(1).

Cohen, Paul R, 1987. "Managing uncertainty", in: *The Third Conference on Artificial Intelligence Applications*, UMassCS, Orlando, Florida.

Cohen, Paul R, 1987. "Steps toward programs that manage uncertainty", in: *Third Workshop on Uncertainty in Artificial Intelligence*, pp. 372–379, American Association for Artificial Intelligence.

Cohen, Paul R and Day, David S, 1988. *The Centrality of Autonomous Agents in Theories of Action Under Uncertainty*, EKSL Technical Report, University of Massachusetts, January 1988. To appear in the *International Journal for Approximate Reasoning.*

Cohen, Paul R, Day, David S, Delisio, Jeff, Greenberg, Michael, Kjeldsen, Rick, Suthers Dan and Berman, Paul, 1987. "Management of uncertainty in medicine", *International Journal of Approximate Reasoning* **1**(1) pp. 103–116.

Cohen, Paul R, Greenberg, Michael and Delisio, Jeff, 1987. "MU: A development environment for prospective reasoning systems, in: *Proceedings of the Sixth National Conference on Artificial Intelligence*, pp. 783–788, Seattle, Washington, July 1987.

Erman, L and Lesser, V R, 1975. "A multi-level organization for problem solving using many diverse, cooperating sources of knowledge", in: *Proceedings of the International Joint Conference on Artificial Intelligence, 1975.*

Erman, L D, Hayes-Roth, F, Lesser, V R and Reddy, D R, 1980. "The HEARSAY-II speech understanding system: Integrating knowledge to resolve uncertainty", *Computing Surveys* **12** pp. 213–253.

Hanks, Steve, 1987. "Temporal reasoning about uncertain worlds", in: *Third Workshop on Uncertainty in Artificial Intelligence*, pp. 114–122, American Association for Artificial Intelligence.

Hayes-Roth, F, 1985. "A blackboard architecture for control", *Artificial Intelligence*, **26** pp. 251–321.

Howe, Adele and Cohen, Paul, 1986. *A Typology for Constructing Decisions*, Department of Computer and Information Science 86-14, University of Massachusetts.

Tong, Richard M and Appelbaum, Lee A, 1987. "Problem structure and evidential reasoning", in: *Third Workshop on Uncertainty in Artificial Intelligence*, pp. 313–319, American Association for Artificial Intelligence.

Wellman, Michael P, 1987. "Qualitative probabilistic networks for planning under uncertainty", in: *Uncertainty in Artificial Intelligence*, John F Lemmer (Ed.).

ON POSSIBILITY THEORY AND INFERENCE MODES – SOME COMMENTS ABOUT "AN AI VIEW OF THE TREATMENT OF UNCERTAINTY" BY A. SAFFIOTTI

Didier Dubois and Henri Prade

Laboratoire Langages et Systèmes Informatiques, Université Paul Sabatier, 118 route de Narbonne, 31062 TOULOUSE, Cédex (FRANCE)

The management of uncertain and incomplete knowledge is an important issue in Artificial Intelligence and there exists an enormous literature which addresses the different facets of the topic. Alessandro Saffiotti provides us with a rich synthetic introduction to most of the proposals on the treatment of uncertainty. We will concentrate our comments mainly along two lines. First, we briefly discuss the classification of existing models for the representation of uncertainty and the place that possibility theory takes among them – an approach not explicitly mentioned by Saffiotti in his review. Then, we shall distinguish between various inference modes which are not specific to a particular uncertainty model.

We have at our disposal the following well-founded numerical models of uncertainty: probability theory, possibility theory, Shafer's theory of evidence, upper and lower probability systems (see Kyburg, (1987) for instance). Possibility theory has been introduced by Zadeh (1978) (see Dubois and Prade (1988) for an introductory book) and enables you to represent states of partial or total ignorance by means of dual possibility and necessity measures Π and N, with $N(A) = 1 - \Pi(\text{not } A), \forall A$; the characteristic axiom is $\Pi(A \cup B) = \max(\Pi(A), \Pi(B)), \forall A, B$. Thus $\Pi(A)$ and $N(A)$ are only weakly related and total ignorance about A is captured by $\Pi(A) = 1$ and $N(A) = 0$. By contrast, probability can model pure randomness $P(A) = P(\text{not } A) = \frac{1}{2}$, (which cannot be done with possibility and necessity, since $\Pi(A) = N(A) = \frac{1}{2}$ is not allowed with ordinary events), but not total ignorance in a satisfactory way. Although a possibility measure can be built from a so-called possibility distribution (which corresponds to the restriction of Π to the singletons), that can be viewed as a fuzzy set, possibility theory and fuzzy sets are two different things which should not be confused. For instance, the intersection of two fuzzy sets F and G is defined in terms of membership functions, following multi-valued logics by $\mu_{F \cap G} = \min(\mu_F, \mu_G)$, while we do *not* have $\Pi(A \cap B) = \min(\Pi(A), \Pi(B))$ in general. Possibility theory provides us with degrees of confidence in the truth or the falsity of a proposition (i.e. to what extent it is possible, that a proposition is definitely true, given the available knowledge which may be incomplete), while fuzzy sets model vague predicates or categories, and induce scalar degrees of truth in the case of complete information (e.g. to what extent it is true that a person, whose height is 1.75 m, is *tall*). Obviously the set-theoretic dimension of fuzzy sets can be conjugated in a nice way with the measure-theoretical aspects of possibility theory thus dealing with uncertain propositions and vague predicates in a common framework. The reader is referred to Dubois and Prade (1988) for further discussions on the differences and complementarities between fuzzy logic and possibilistic logic. Many different interpretations of the same mathematical framework also exist (see Fine (1973) for the probabilities). Thus, probability measures on the one hand and possibility and necessity measures on the other hand can be considered, mathematically speaking, as extreme particular cases of Shafer's plausibility and belief (support) functions (see Dubois and Prade (1982) for instance). Plausibility and belief functions, as well as possibility and necessity measures may also be viewed as to particular cases of upper and lower probability systems (possibility theory is the simplest of such systems). However possibility theory can also be used without any reference to probability. Shafer's theory of evidence can also be thought of in terms of random sets (see Dubois and Prade (1986) for such a presentation), thus leading to a set-theoretic view of this approach. Another important representation issue connected with numerical models and particularly with probability theory, is the choice between material implication and conditioning when modelling "if... then" rules. See Calabrese (1987) for instance, for a discussion.

Apart from well-founded numerical models, there exist on the one hand empirical models, such as the MYCIN calculus of uncertainty or Cohen's endorsements, and on the other hand purely symbolic approaches, based on non-standard logics. In both types of framework valuable ideas are introduced and implemented. Reiter's default logic (1980) is particularly worth mentioning as an example of the second type of approach. The reader interested in a comparison of symbolic and numerical approaches on the same example with a discussion of the representation capabilities and the non-monotonic behaviours of each formalism, is referred to the recent study by the group Léa Sombé (1988).

There are two general inference mechanisms, which are not specific to a particular numerical model, and which must be distinguished. The first one consists of assuming the existence of a unique joint confidence measure which is constrained by the different pieces of available information; Nilsson's probabilistic logic, Quinlan's INFERNO's system (1983), Dubois and Prade's resolution method (1987) in possibilistic logic are examples of this approach. Confidence in propositions of interest, induced by the constraints on the confidence measure, can then be computed. A second general method combines all the pieces of information, each of them being represented in terms of a single (joint) distribution, and projects the resulting distribution on the domains of variables in which we are interested; Zadeh's approximate reasoning techniques (1979), Dempster's rule of combination, or Pearl's implementation of Bayesian inference (1986) are examples of this second approach. These two kinds of approaches can be used either at a local level (e.g. as in production rule systems where rules are triggered separately and partial conclusions are then combined which may lead to dubious results in some cases) or at a global level taking into account the whole knowledge base (this guarantees the validity of the results but may be computationally untractable). In order to make a compromise between the preference for local computations and the validity of the results, we may either introduce stochastic independence hypotheses as in Pearl (1986) or decompose the knowledge base into logically independent sub-bases (Chatalic et al., 1987).

Note: The references used in this discussion, which can already be found in Saffiotti's paper are not repeated in the following list.

References

Léa Sombé (Besnard, P, Cordier, M O, Dubois, D, Fariñas del Cerro, L, Froidevaux, C, Moinard, Y, Prade, H, Schwind, C and Siegel, P) (1988). "Inférences non-classiques en intelligence artificielle-Ebauche de comparaisons sur un exemple", *Actes 2èmes Journées Nationales du P.R.C.-G.R.E.C.O. Intelligence Artificielle*, Toulouse, 14–15 Mars, 1988, Teknéa, Marseille, pp. 137–230.
Calabrese, P, 1987. "An algebraic synthesis of the foundations of logic and probability", *Information Sciences* **42** pp. 187–237.
Chatalic, P, Dubois, D and Prade, H, 1987. "A system for handling relational dependencies in approximate reasoning", *Proc. 3rd Inter. Conf. on Expert Systems*, London, June 2–4, 1987, 495–502.
Dubois, D and Prade, H, 1982. "On several representations of an uncertain body of evidence", in: *Fuzzy Information and Decision Processes*, M M Gupta and E Sanchez, (Eds.), North-Holland, pp. 167–181.
Dubois, D and Prade, H, 1988. *Théorie des Possibilités. Applications à la Représentation des Connaissances en Informatique*, Masson, Paris, 1985, (2nd revised and augmented edition, 1987). English translation: *Possibility Theory. An Approach to the Computerized Processing of Uncertainty*, Plenum Press, New York, 1988.
Dubois, D and Prade, H, 1986. "A set-theoretic view of belief functions: logical operations and approximation by fuzzy sets", *Inter. J. General Systems* **12** pp. 193–226.
Dubois, D and Prade, H, 1987. "Necessity measures and the resolution principle". *IEEE Trans. on Systems, Man and Cybernetics* **17** pp. 474–478.
Dubois, D and Prade, H, 1988. "An introduction to possibilistic and fuzzy logics (with discussions and a reply", in: *Non-Standard Logics for Automated Reasoning* (E H Mamdani, P Smets, D Dubois, H Prade, (Eds.), Academic Press, New York, pp. 287–326.
Fine, T, 1973. *Theories of Probability*, Academic Press, New York.
Kyburg, H, 1987. "Bayesian and non-Bayesian evidential updating", *Artificial Intelligence* **31** pp. 271–293.
Pearl, J, 1986. "Fusion, propagation and structuring in belief networks", *Artificial Intelligence* **29** pp. 241–288.
Quinlan, J R, 1983. "INFERNO: a cautious approach to uncertain inference", *The Computer Journal* **26** pp. 255–269.
Reiter, R, 1980. "A logic of default reasoning", *Artificial Intelligence* **13** pp. 81–132.
Zadeh, L A, 1978. "Fuzzy sets as a basis for a theory of possibility", *Fuzzy Sets and Systems* **1** pp. 3–28.
Zadeh, L A, 1979. "A theory of approximate reasoning", in: *Machine Intelligence* (J E Hayes, D Michie, L I Mikulich, (Eds.), Elsevier, NY Volume **9** pp. 149–194.

A SYSTEM PERSPECTIVE FOR CHOOSING BETWEEN DEMPSTER-SHAFER AND BAYES

John F Lemmer

Knowledge Systems Concepts Inc, Rome, NY 13440, USA

The purpose of this discussion is to point out the double standard which persists when Bayesian and Dempster-Shafer approaches to the management of uncertainty are compared. The discussion presupposes acceptance of at least a portion of the eclectic position: that different representations may be needed on different occasions. We will not argue that either approach is always better. Instead what follows is an attempt to compare even-handedly the advantages and drawbacks of the Bayesian and Dempster-Shafer approaches within a particular context: building a state of the art AI system which can manage uncertainty. From an engineering perspective, this discussion can be viewed as providing data to support a tradeoff analysis between Bayesian and Dempster-Shafer approaches but without providing the evaluation criteria.

The data provided here amplifies issues raised in Saffiotti's article, either in the discussion of Bayesian or DS techniques. Saffiotti raises most of the interesting issues and is more even-handed than most in his comparison of these two techniques. But as we shall see, he too has fallen somewhat into the mode of pointing out Bayesian weaknesses while extolling Dempster-Shafer virtues.

There are two major dimensions that an even-handed tradeoff analysis of Bayesian and Dempster-Shafer techniques ought to address. The first dimension is cost-effectiveness: How much better do DS techniques handle "certainty about certainty" and is this improvement worth the increased computational and human resources required by DS when compared to Bayes. Saffiotti alludes to the computational costs but fails to mention the human costs, costs primarily associated with the knowledge engineering of belief functions. The second dimension is applicability of the uncertainty calculus: is Bayes' rule or Dempster's rule more applicable to the system goals, the problem domain, and the type of knowledge available.

Cost-effectiveness

There is no doubt that Dempster-Shafer theory currently provides much more natural and much more developed techniques for providing certainty about certainty. Even though, as Saffiotti mentions, Bayesian advocates claim that, *in principle*, Bayesian techniques can provide certainty about certainty, this is not any different than claiming that in principle, a Turing Machine can do anything a LISP machine can do. At the present time, the *technology* of Bayesian certainty about certainty is undeveloped, and cannot be readily compared with the DS approach.

There is no dispute that Bayesian approaches offer a well-developed decision theory and DS approaches do not. Work is proceeding on a decision theory from probability intervals. But we are unaware of the development of any decision theories based directly on belief functions. Thus DS decision support is like Bayesian certainty about certainty: not yet ready for application.

What price should an AI system be willing to pay for certainty about certainty? There is no doubt that such information is desirable to support meta-reasoning. Some may even regard such information as a *sine qua non*. But at the current time, the *value* of certainty about certainty in the sense of improved system performance and acceptability is only a matter of conjecture. Indeed, at the present time, lack of a decision theory which utilizes such information makes anything more than conjecture about the value very difficult.

How important is a decision theory for an AI system? This depends on the goals of the system. If any automatic decisions are to be made, even at a low level, it would seem to be an important consideration. If the use of the system is to support a human decision maker, a decision theory may not be required.

Two conclusions can be drawn from this. The first is that advocates of Bayesian approaches to certainty about certainty must develop a technology supporting this approach if they expect it to be incorporated into AI systems requiring such information. The second is that the cost of certainty about certainty must be weighed against the value of a decision theory in any system design. Even if DS techniques will add value to the overall system, they are extremely costly in comparison to Bayesian techniques. The computation and knowledge engineering costs of Dempster-Shafer are intrinsically much higher than costs of Bayes, higher by an exponential factor. This cost difference arises directly from differences in information representation in the two schemes. For small problems the cost difference may well be worthwhile, but for large problems the cost difference may be prohibitive. The comparison below is based on the "full theories". Less costly techniques approximating the full theories have been developed for both Bayesian and DS, as noted by Saffiotti. But we feel that a comparison of the full theories is still meaningful because it is difficult to imagine how the intrinsic complexity differences can be overcome without "throwing the baby out with the bath water" while simplifying DS. Saffiotti claims that "the main problem of the Bayesian approach . . . is that the method requires a bulk of data (*a priori* and conditional probabilities) which is often hard – if not impossible – to come by".

Let us compare the bulk of data required by Bayesian representation with that required by Dempster-Shafer representation. A set of n mutually exclusive and collectively exhaustive (probabilistic) events is analogous to a DS frame of discernment containing n elements. It takes n numbers to represent a probability distribution over such a set of events. It takes 2^n numbers to represent a belief function over the equivalent frame of discernment. If $n = 10$, it takes 10 numbers to represent the probability distribution and 1024 numbers to represent the belief function.

In most practical problems the events of interest will not be mutually exclusive. Suppose that instead of 10 discrete events, the nature of the domain is such that there are 10 binary (potentially) overlapping events. Then the (Bayesian) representation of the probability distribution requires 1024 numbers. But since the frame of discernment must contain 1024 elements,[1] the full belief function requires 2^{1024} numbers, that is, more than 10^{308} parameters!

Where are all of these parameters to come from? It is certain that in neither case will they be elicited from experts directly. The answer for Bayesian techniques has been to elicit a small number of related parameters which under certain assumptions implicitly specify the prior. Saffiotti's equation (4) is an example of such a method. The work of Pearl (1986), Cooper (1984), and Spiegelhalter (1986) can all be interpreted in this light.

Where the parameters come from for DS techniques is not clear. Most published examples of Dempster-Shafer go directly from a verbal description of the evidence to a belief function, implying that the elements of the belief function are elicited from experts. We are not aware of any general techniques yet developed for eliciting belief functions. It is well-known that eliciting probabilities is fraught with difficulties. Eliciting belief functions should be much more difficult still, not only because there are so many more, but also because their meaning is even less intuitive than the meaning of probabilities.

Another approach implicit in many discussions of DS application is to elicit probability intervals for each element of the frame of discernment. From these intervals implicitly or explicitly generate a belief function satisfying the interval constraints. (We take implicit generation of a belief function to include manipulation of the intervals themselves.) If explicit generation is attempted, computational problems remain. If we are to place any reliance on the result of direct manipulation of intervals, we must understand the assumptions justifying such manipulations in terms of DS theory. Surface level analogies do not provide such understanding as is clear from the history of MYCIN certainty factors. MYCIN certainty factors, now rejected by the very group which developed them (Heckerman, 1986), are an example of techniques accepted by many because of their surface similarity to Bayesian techniques. Their full analysis in light of Bayesian theory revealed a number of unwarranted hidden assumptions. One should be wary of ad hoc techniques which allege to be DS based.

Even if intervals should prove to adequately support reasonable approximations of DS, several significant knowledge engineering problems remain. First, not all sets of probability intervals on the frame of discernment are feasible.[2] Secondly, probability intervals on elements of the frame of discernment severely underconstrain the belief function. How is the particular belief-function for (implicit or explicit) updating to be chosen from the large set of feasible solutions? (Similar problems of underconstraint in Bayesian approaches are usually settled by invoking independence or maximum entropy arguments.) What are the corresponding assumptions in DS theory? How acceptable are these assumptions? How can we assess whether or not they apply to the problem domain?

A fair comparison of knowledge engineering to support Bayesian and DS techniques, might be that while there are known serious hazards on the sea of Bayesian knowledge engineering, the sea of DS knowledge engineering is virtually uncharted.

Not only is the DS sea uncharted, but consider the following. It seems from the above, that generation of a belief function concerning a body of evidence is at least as difficult and probably more so, than generating a Bayesian prior (a known hard problem). Not only that, but more than one belief function must be generated; one must be generated for each evidential source!

Bayes' rule vs Dempster's rule

Bayes' rule and Dempster's rule seem to be applicable to different types of evidence. For Bayes' rule to be appropriate, there must be adequate knowledge of a prior distribution. For Dempster's to be appropriate, the bodies of evidence to be combined must be "independent". There is little

[1] The elements of a frame of discernment must be mutually exclusive as must the joint events of a probability distribution.

[2] As simple examples, there is no belief function over $\{a, b\}$ satisfying $.4 \leqslant p(a) \leqslant .6$, and $.5 \leqslant p(b) \leqslant .6$, nor over $\{a, b, c\}$ satisfying $.45 < a < .46$, $.30 < b < .315$, and $.20 < c < .23$, nor even satisfying $x < a < x + .01$, $y < b < y + .015$, and $z < c < z + .03$ for any non-negative x, y, and z. Nonetheless there are probability distributions which do satisfy these constraints.

disagreement on either of these points. But there is confusion on how knowable priors can be and what it means for bodies of evidence to be independent. We will examine each of these issues below. But first we clarify the relation between Bayes' rule and Dempster's rule.

It is a widely held belief that Bayes' rule is a special case of Dempster's rule. It is true that if two belief functions are such that events possible in one are ruled out in the other, Dempster's rule will give the same result as Bayes' rule. This occurs when the less focused belief function is treated as the Bayesian prior and the updating event is the (now) certain event (in terms of the second belief function) that none of the ruled out events have occurred.[1] In other cases, the rule provides different answers and, as we will see below, the Bayesian answer is sometimes to be preferred.

Acquisition of a meaningful prior distribution is indeed a difficult task. In the past Bayesians have often sidestepped this problem by assuming priors which are difficult to justify. Often the principle of maximum entropy has been invoked, indeed so often that the notion that "the probability of everything is equal, *a priori*" was incorrectly viewed by some as an intrinsic component of the Bayesian methodology.

The problem of priors can no longer be used *ipso facto* to rule out the Bayesian approach on epistemological grounds. The problem of generating meaningful priors is being successfully addressed in a number of areas. For example the current work in causal modelling can be viewed as developing techniques for generating priors from causal knowledge. Other work is aimed towards generating priors from large bodies of statistical information.

The applicability of Dempster's rule seems to be a cloudy issue at the current time because the concept of "independent bodies of evidence" does not seem to be clearly defined. However, it is possible to demonstrate by example that certain bodies of evidence should *not* be combined with Dempster's rule. In (Lemmer, 1986) it was shown that evidence from independently operated but low error rate sensors[2] should not be so combined in order to draw inferences about a population. The basic and paradoxical problem encountered was that the low error rate of the sensors rendered their evidence mutually dependent.

In (Lemmer, 1986) it was further suggested, based on attempts at the construction of hypothetical sample spaces, that DS techniques are most appropriate when there is no knowledge concerning the sample space other than that contained in the evidence. It appears that Dempster's rule can be best understood as a method for *constructing* a sample space rather than for encoding or learning information about a sample space which, conceptually at least, actually exists.

Thus we can interpret Saffiotti's comparison which he claims "may suggest the superiority of Dempster-Shafer theory" as meaning that DS theory seems to offer a number of theoretical advantages if its practical difficulties can be overcome and if the use of Dempster's rule is appropriate. To us, DS theory seems to be rich with research opportunities, but not yet ready for full scale application. We feel that it is too early to know if it will ever be practical in any but small problems. We suspect that the current absence of a long list of known problems with DS, similar to lists of Bayesian problems, is due to the current level of application of DS: small, technology driven applications designed to showcase the advantages of DS. We predict that as larger, result driven applications are attempted, the list of DS problems will become longer!

References

Cooper, G F, 1984. *NESTOR: A Computer-based Medical Diagnostic Aid that Integrates Causal and Probabilistic Knowledge*, PhD, Th., Computer Science Department, Stanford University, Rep. No. STAN-CS-84-48.

Heckerman, D, 1986. "Probabilistic interpretations for MYCIN's certainty factors", in: *Uncertainty in Artificial Intelligence*, Kanal and Lemmer (Eds.), North-Holland.

[1] Shafer, on page 66 of (Shafer, 1976) discusses this in terms of the disjunction of a set events being ruled in with certainty. But the essential relation to Bayes is that the disjunction of some other set of events has been ruled out with certainty. It should be noted that although Baye's rule can be applied in this situation, its real power is achieved when a conjunction of events is ruled in with certainty and is desired to estimate the probabilities of this conjunction when extended to include other events.

[2] E.g. radar and infra-red sensors.

Lemmer, J F, 1986. "Confidence factors, empiricism and the Dempster-Shafer theory of evidence", in: *Uncertainty in Artificial Intelligence*, Kanal and Lemmer (Eds.), North-Holland.
Pearl, J, 1986. "A constraint-propagation approach to probabilistic reasoning", in: *Uncertainty in Artificial Intelligence*, Kanal and Lemmer (Eds.), North-Holland.
Shafer, G. 1976. *A Mathematical Theory of Evidence*, Princeton University Press.
Spiegelhalter, D J, 1986. "Probabilistic reasoning in predictive expert systems", in: *Uncertainty in Artificial Intelligence*, Kanal and Lemmer (Eds.), North-Holland.

COMMENTARY ON "AN AI VIEW OF THE TREATMENT OF UNCERTAINTY" BY ALESSANDRO SAFFIOTTI

David J. Spiegelhalter
MRC Biostatistics Unit, Cambridge, CB2 2BW

The ecumenical spirit of this paper should be welcomed by all those seeking procedures for handling uncertainty that are both theoretically sound and pragmatically appropriate. My brief comments will concern the role of rigorous probabilistic reasoning in expert systems, as opposed to AI in general.

Saffiotti has made a number of important points concerning probabilistic reasoning that are very positive for someone who is not, apparently, a card-carrying committed Bayesian. However, I feel I must first point out that the description of the use of Bayes theorem in AI (section 2.2), while no doubt descriptive of much of current practice, does not reflect recent developments. There has been an exciting growth of practical procedures for handling probabilistic relationships on unknown quantities whose interrelationships may be expressed graphically (see, for example, Pearl, (1986) and Lauritizen and Spiegelhalter (1988)) and the rigorous techniques supersede previous *ad hoc* approaches in MYCIN or PROSPECTOR. In this representation, "extending the conversation" (section 2.5) is certainly not a "hornet's nest", since additional events upon which conditioning is of interest are simply represented as additional nodes in the graph. In particular, the conditional probabilities themselves may be considered as unknown quantities (Spiegelhalter, 1986) and distributions over these reflect the vagueness or imprecision that is inevitable whether assessments are based on subjective opinion or statistical data. (Incidentally, the vignette of section 4.1 reveals the absurdity of trying to assess quantities such as MYCIN's certainty factors which supposedly combine both belief and importance, but does not in any way invalidate the subjective assesment of well-defined conditional probabilities.) As Saffiotti very nicely observes in section 2.5, explicit recognition of imprecision in numbers shifts the problem from an epistemological to a computational level, but this is fine since the computational problems are no longer so daunting and, if necessary, can always be solved using stochastic simulation techniques. Saffiotti also notes in section 3.4 that the vagueness of a probability value and its sensitivity to further experience are two distinct concepts, but while this may be true from an epistemological point of view, they are identical computationally given the representation described above: "vagueness" of a specific conditional probability p can be assessed by searching over imprecise probability values that influence p in order to generate a distribution over p, while "sensitivity" of p is obtained by carrying out the same searching and averaging procedure with respect to a different set of nodes – those that are not yet currently observed on the case in hand, but are potentially so.

I therefore strongly agree with Saffiotti's perceptive point (section 4.2) that the certainty itself becomes an object to be manipulated. Within a probabilistic framework this is quite feasible; as discussed above, the probability tables themselves become objects with their *own* uncertainty, which can themselves be manipulated according to the laws of probability, since, as Pearl (1987) emphasises, the probability calculus operates identically at all levels. In particular, the probability calculus is used not only to manipulate evidence concerning an individual case, but also to revise opinion concerning the probabilities as experience grows from case to case. This ability to operate at both levels is extraordinarily powerful.

I did not, however, mean this commentary to be another ideological polemic, and I do believe that probability is not appropriate for representing all aspects of what is linguistically covered by the term "uncertainty". At a minimum, it will be very interesting in the future to investigate the idea of handling both qualitative and quantitative descriptions within the same basic representation. It should then be possible to simultaneously reason numerically and to justify the conclusion in qualitative terms with regard, say, to the source and precision, or even to argue qualitatively as a first pass or to structure the problem. I would, though, take issue with Saffiotti's apparent identification of qualitative with "explicit" and "quantitative" with "implicit". What could be more explicit than a number together with a representation of its source?

Whether a qualitative argument is sufficient depends on the context and objectives. This is where I feel it is important to distinguish between AI ("science") and expert systems ("technology"); I have never really been clear about the objectives of the former, but the latter must surely be evaluated, at least partly, on the basis of its ability to cope with repeated use in a similar context, since by definition an expert system will be handling data on many cases. I would, while trying not to be too dogmatic, claim that the probability calculus is the only mechanism for dealing with systematic criticism and improvement, as case data accumulate, of both the structure of a system *and* its attendant uncertainty. As expert systems become more accepted, the recognition of this claim will, I believe, lead to a radical shift in the argument concerning uncertainty in expert systems.

References

Lauritzen, S L and Spiegelhalter, D J, 1988. "Local computations with probabilities on graphical structures and their application to expert systems (with discussion)". *J. Roy. Statist. Soc.*, **B**, 50 (to appear).

Spiegelhalter, D L, 1986. "Probabilistic reasoning in predictive expert systems", in: *Uncertainty in Artificial Intelligence*, Kanal, L N and Lemmer, J (Eds.), North Holland, Amsterdam, pp. 47–68.

Pearl, J, 1987. "Do we need higher-order probabilities and, if so, what do they mean?" *Proceedings of Third AAAI Workshop on Uncertainty in Artificial Intelligence*, Seattle, pp. 47–60.

ANSWER TO "AN AI VIEW OF THE TREATMENT OF UNCERTAINTY" BY ALESSANDRO SAFFIOTTI

Philippe Smets

IRIDIA, Université Libre de Bruxelles, 50 av. F. Roosevelt, CP 194/6, 1050 Bruxelles, Belgium

The author must be congratulated for his open minded presentation of the various models proposed to represent and manage uncertainty. His presentation of Dempster-Shafer theory corresponds to the one usually encountered today in the AI community. We do not share this interpretation and summarize the one we derived from Shafer's original work as presented in his book *A mathematical theory of evidence*.

Saffiotti claims that "the probability of A is bounded within the interval [spt(A), pls(A)]", implying the existence of a single valued probability p(A). So Dempster-Shafer theory as presented is a generalization of Bayesian theory, where each proposition A receives a single valued probability p(A) (an additive measure of course) but where we only know that p(A) is within a certain interval. This is a special form of upper and lower probabilities theory initially developed in the early fifties by I. Good (see Good, 1983). Then the relation of these probabilities with non-negative "probability masses" and the use of Dempster's rule of combination for probability kinematics, the two major ingredients of the model, are essentially *ad hoc*, even though natural.

1 The transferable belief model

Our interpretation of Shafer's early work is different. We do not consider the notion of additive probability. We start by postulating that the impact of a piece of evidence on a finite frame of discernment Θ consists in allocating parts of an initial unitary amount of belief among the subsets of Θ. For $A \subseteq \Theta$, m(A) is a part of our belief that supports A and that, due to lack of further information, does not support any strict subset of A. We call m(A) a basic belief mass.

To get a clear understanding of the model, consider a murder case with three suspects: Robert, John and Mary, i.e. $\Theta = \{R, J, M\}$. Suppose a witness tells us that he thinks he saw the killer and that he saw a male. Hence his testimony supports $R \cup J$. But he might be wrong. Suppose we believe at level α that he indeed saw the killer. Hence $m(R \cup J) = \alpha$ and $m(R \cup J \cup M) = 1 - \alpha$. The part α of our initial belief is given to $R \cup J$, but cannot be given more specifically to either R of J given the lack of further information. The usual presentation of Dempster-Shafer theory would require the existence fo some β value such that $P(J) = \beta$ and $P(R) = \alpha - \beta$ but such that *we* only know that β belongs to $[0, \alpha]$. This requirement is unnecessary and could even lead to paradoxical conclusions because postulating the existence of this (numerically unknown) β implies evidently some extraneous constraints on the models.

Suppose then that we receive new evidence that assures us that the truth is in $\Theta' \subseteq \Theta$. Then a basic belief mass given to $A \subseteq \Theta$ will now support $A \cap \Theta'$. If we learn in the murder case that John was not the killer, then the testimony about gender supports Robert. Hence if m' are the basic belief masses after the "conditioning" evidence, then $m'(R) = \alpha$, $m'(R \cup M) = 1 - \alpha$, and all other m' are null. How to handle masses m' given to $\emptyset$ is discussed at length in Smets (1988a) where open and closed worlds assumptions are defined. This belief kinematic corresponds, in fact, to the so-called Dempster's rule of conditioning. It satisfies the nature of the basic belief masses that are allocated to some set A but that are free to be transferred to any subset of A should new evidence become available. Hence the name of *transferable belief model*. It is important to realize that this belief kinematic is not just an *ad hoc* rule but corresponds to the intrinsic nature of the basic belief masses model.

A set of natural requirements suffices then to justify Dempster's rule of combination as the unique rule for combining basic belief masses derived from two distinct pieces of evidence (Smets, 1988b).

Given the potentially transferable basic belief masses, one defines the degree of belief (plausibility) of A as the sum of all the basic belief masses that support necessarily (possibly) A, i.e. the basic belief masses allocated to sets B that are non-empty subsets of (compatible with) A.

$$\mathrm{bel}(A) = \sum_{X \subseteq A X \neq \emptyset} m(X)$$

$$\mathrm{pl}(A) = \sum_{X \cap A \neq \emptyset} m(X)$$

(bel and pl correspond to spt and pls in Saffiotti's paper).

2 Decision mechanism

The design of the decision mechanism is not clear within the Dempster-Shafer theory when it is interpreted as an upper and lower probabilities theory. But within the transferable belief model approach, it is perfectly clear and an unique solution can be derived. One must realize that the transferable belief model does describe in fact our *credal state*, i.e. our subjective and personal judgement about which subset $A \subseteq \Theta$ contains the truth. When we must decide (or bet), we must construct a pignistic probability on Θ (pignus = a bet in Latin). Indeed all betting behaviours should be organized according to a probability function if we want to avoid Dutch book arguments, etc.... Such probability must of course be based on our credal state on Θ, i.e. on those basic belief masses we give to Θ.

If one postulates that the pignistic probability BetP(A) given to A does not depend on how the basic belief masses are split among the subsets of $\bar{A}$, as far as their total bel($\bar{A}$) remains constant, then the unique solution for Bet $P:\Theta \rightarrow [0, 1]$ is

$$\mathrm{BetP}(A) = \sum_{X \subseteq \Theta} m(X)\,|A \cap X|/|X|$$

BetP is indeed a probability function (proofs in Smets, 1988c).

This solution depends on what are the singletons of Θ, i.e. on the betting framework. Its structure and impact on bets are described in detail in Smets (1988c).

In conclusion, we have given a short description of the interpretation we derived from Shafer's initial work as described in his book, not as described in his later papers where he moves toward the Dempster-Shafer model, i.e. a restricted upper and lower probabilities model enhanced by Dempster's rule of combination but apparently poised by ad hockeries. Our transferable belief model is completely free from any probabilistic connotation and permits the construction of a model that keeps all the power and beauty of Shafer's model while hopefully avoiding its criticisms.

References

Good, I J, 1983. *Good thinking. The foundation of probability and its applications.* Univ. Minnesota Press, Minneapolis, Min.

Smets, Ph, 1988a. "Belief functions", in: Smets, Ph, Mamdani, A, Dubois, D and Prade, H (Eds.). *Non standard logics for automated reasoning*, Academic Press, London pp. 253–286.

Smets, Ph, 1988b. "The combination of evidence in the transferable belief model" (submitted for publication).

Smets, Ph, 1988c. "The transferable belief model: comparison with Bayesian models", (submitted for publication).

ON THE TREATMENT OF UNCERTAINTY IN AI

L. A. Zadeh*

In *An AI View of the Treatment of Uncertainty*, Dr Saffiotti presents an insightful analysis of some of the complex issues which relate to the management of uncertainty in AI. Generally, I agree with much of what Dr Saffiotti has to say. There are, however, some important issues which I see in a different light.

First, there is a widespread and, in my opinion, misplaced belief that there exists a collection of *n* distinct techniques for dealing with uncertainty which can be ordered linearly in terms of their effectiveness. The point I am trying to make is that in most cases the techniques in question cannot be ordered linearly because they have different domains of applicability, with the understanding that applicability, or effectiveness, is a matter of degree. This is a consequence of the fact that uncertainty is a concept with many dimensions, and that a technique which is effective in relation to a particular dimension may be lacking in effectiveness along others.

In more specific terms, the principal facets of uncertainty which are of relevance to AI are the following.

1 *Lexical imprecision or, equivalently, lexical elasticity.* For example, in the proposition

"The price of oil is not *likely* to increase *substantially* in the *near* future"

the italicized words are lexically imprecise in the sense that the predicates which they represent are fuzzy rather than crisp. Lexical imprecision may be viewed as a concomitant of elasticity of meaning.

* Computer Science Division, University of California, Berkeley, CA 94720. Research supported by NASA Grant NCC-2-275 and NSF Grant DCR-85-13139.

2 *Granularity of information.* Granularity relates to the incompleteness of information regarding the values of a variable. For example, the proposition "Mary was born in March" conveys incomplete information regarding the birthdate of Mary, with March playing the role of a crisp granule. More generally, granularity relates to propositions of the form "X is A" where X is a variable taking values in a universe of discourse U, and A is a crisp or fuzzy subset of U. For example, in the case of the proposition "Mary is young", X is the age of Mary and *young* plays the role of a fuzzy granule which constrains X. In fuzzy logic, A defines the *possibility distribution of* X, i.e. the fuzzy set of values which X can take (Zadeh, 1979; Zadeh, 1981; Zadeh, 1983).

3 *Probability assignment and probability qualification.* Probability assignment has to do with the specification of the probability distribution of a variable. In symbols, it may be represented as the proposition X *isp* P, where P is the probability distribution of X and the copula *isp* serves to indicate that what is assigned to X is its probability distribution, P. By contrast, the copula *is* in "X is small", indicates that what is assigned to X is its possibility distribution, *small.*

Probability qualification has to do with the assignment of a probability to a proposition or, equivalently, to an event. For example, the proposition, "The probability that Mary was born in March is 0.8" is a probability qualification of the proposition "Mary was born in March". In the case of the probability-qualified proposition "It is likely that Mary is young" the proposition "Mary is young" is fuzzy and so is the qualifying probability *likely.*

The three facets described above are not independent and may occur in various guises and different combinations. Nevertheless, it is convenient to assess the effectiveness of various methods of dealing with uncertainty in terms of their capability to deal with the problem of inference from premises which exhibit the facets in question singly or in combination.

In a general way, Bayesian methods naturally fit those situations in which (a) the premises involve probability assignments; (b) there is no lexical imprecision; and (c) there is no granularity. In the same frame of reference, the Dempster-Shafer theory is effective when (a) the premises involve probability qualification; (b) there is no lexical imprecision; and (c) certain restricted types of granularity are allowed. Viewed in this perspective, the domain of effectiveness of the Dempster-Shafer theory is broader than that of the Bayesian methods.

What is important to note is that neither the Bayesian methods nor the Dempster-Shafer theory in its standard form address the issue of lexical imprecision. To deal with this facet of uncertainty, it is natural to employ fuzzy logic (Zadeh, 1983).

As an illustration, consider the following problems.

(a) A, B and C are three points in a plane. The distance between A and B is approximately a; the distance between B and C is approximately b. What is the distance between A and C?
(b) X is a large number; Y is much larger than X. What can be said about the sum of X and Y?
(c) There are twenty tenants in an apartment house. A few are very young; several are young; a few are old; and the rest are middle aged. How many are not very young and not very old?
(d) X, Y and Z are three variables whose relationship is described by the following rules:

Z is small if X is large and Y is medium
Z is medium if X is small and Y is large
Z is large if X is medium and Y is small

The question is: What is the value of Z when X is not small and Y is medium?

In the case of (a), it is simplest and most natural to use possibility theory, which is a branch of fuzzy logic (Zadeh, 1981), or, equivalently, a version of fuzzy Prolog (Baldwin, 1987; Baldwin et al., 1987; Hinde, 1986; Kanai and Ishizuka, 1986; Mukaidano, Shen and Ding, 1987). One could employ the Bayesian approach if *approximately a* and *approximately b* are interpreted as probability distributions. However, the analysis would be complex and unnatural. There would be no point in using the Dempster-Shafer theory or the theory of endorsements.

In the case of (b), it is again simplest and most natural to employ possibility theory or fuzzy Prolog, since Prolog is basically a possibilistic language. As for the Bayesian approach, a problem which arises is that the predicate *large number* cannot be interpreted as a probability distribution on the real line. However, it can be interpreted as the conditional probability of *large number* given X, or, equivalently, as a random set (Goodman and Nguyen, 1985; Orlov, 1980). These interpretations will lead to the same result as that yielded by the application of possibility theory, but in a much more convoluted way.

In the case of (c), the natural thing to do is to employ an extension of the Dempster-Shafer theory to fuzzy basic probability numbers and fuzzy focal sets (Dempster, 1967; Shafer, 1976). To employ Bayesian methods would be neither simple nor natural.

In the case of (d), if the predicates *small*, *large*, and *medium* are interpreted as possibility distributions, the answer can readily be computed through an application of the rules of inference of fuzzy logic (Zadeh, 1983). Such rules of inference are embodied in the Togai-Watanabe fuzzy logic chip and Yamakawa's fuzzy computer (Preprints, 1987; Togai and Watanabe, 1973; Watanabe and Dettloff, 1987), and have formed the basis for industrial applications of fuzzy control (Sugeno, 1985).

By contrast, it would be hard to solve the problem in question by the use of Bayesian or Dempster-Shafer methods. What this demonstrates once again is that the latter techniques are not well-suited for dealing with problems in which lexical imprecision plays an important role.

It should be noted that problem (a) points to a significant limitation of the conventional methods of dealing with uncertainty in AI. More specifically, one of the most basic rules of inference in evidential reasoning involves the use of *modus ponens*. However, there are many cases, exemplified by problem (d), where there is no rule in the knowledge-base whose antecedent matches exactly a given fact or the consequent of another rule. In this event, neither *modus ponens* nor Bayesian methods are applicable. By contrast, in fuzzy logic the problem can be handled simply and directly through the use of the generalized *modus ponens*, which may be expressed as the rule

$$\begin{array}{l} Y \text{ is } B \text{ if } X \text{ is } A \\ \underline{X \text{ is } A'} \\ Y \text{ is } B' \end{array}$$

where A, B, A' and B' are fuzzy predicates, with A' not required to be equal to A. In terms of A, B and A', the membership function of the fuzzy predicate B' is given by

$$\mu_{B'}(v) = sup_u(\mu_{A'}(u) \bigwedge (1 - \mu_A + \mu_B(v))$$

where μ_A, μ_B, $\mu_{A'}$ and $\mu_{B'}$ are the membership functions of A, B, A' and B', respectively, $\bigwedge = min$, and the supremum is taken over the domain of u. For example

$$\begin{array}{l} Y \text{ is large if } X \text{ is small} \\ \underline{X \text{ is medium}} \\ \quad Y \text{ is } R \end{array}$$

where

$$\mu_R(v) = sup_u(m_{medium}(u) \bigwedge (1 - \mu_{small}(u) + \mu_{large}(u))$$

What emerges from this discussion is that neither bivalent logic nor the Bayesian methods nor the Dempster-Shafer theory provide a natural way of dealing with the basic problem of chaining when there is a mismatch in the bridging links.

Another important point which should be borne in mind when we attempt to compare various methods is that the mere fact that method A is, in principle, a special case of method B does not imply that A is subsumed by B. For example, multivalued logic is a special case of fuzzy logic, but the agenda of multivalued logic and its aims are quite different from those of fuzzy logic. Another example is the relation between linear system theory and nonlinear system theory. Although the latter is more general, its primary concern is with problems which fall outside the domain of applicability of linear system theory. The same applies to the relation between fuzzy logic, the Bayesian methods and the Dempster-Shafer theory, each of which has its own agenda and domain of applicability, regardless of the arguments concerning the issue of which is a special case of what.

In conclusion, Dr Saffiotti deserves to be complimented for presenting an informative analysis of some of the complex issues which arise in the treatment of uncertainty in AI. However, an important issue which is not discussed adequately in Dr Saffiotti's paper is that of lexical imprecision. It is this basic facet of uncertainty that is not addressed by the techniques employing classical probability theory and bivalent logic.

References and related publications

Baldwin, J F, 1987. *SLOP—A System for Support Logic Programming*, ITRC Research Report, University of Bristol, England.

Baldwin, J F, Martin, T P and Pilswood, B W, 1987. "The implementation of FPROLOG—A Fuzzy Prolog interpreter", *Fuzzy Sets and Systems* **23** 119–129.

Bonissone, P P, 1983. "A survey of uncertainty representation in expert systems", *Proc. Second Workshop of the North-American Fuzzy Information Processing Society*, Schenectady, N.Y., General Electric Corporate Research and Development.

Cheeseman, P, 1985. "In defence of probability", *Proc. Int Joint Conference on Artificial Intelligence* Los Angeles, 1002–1009.

Cohen, P, 1985. *Heuristic Reasoning about Uncertainty: An Artificial Intelligence Approach*, Boston: Pitman.

Dubois, D and Prade, H, 1985. *Théorie des Possibilités. Applications a la Representation des Connaissances en Informatique*, Paris: Masson.

Dempster, A, 1967. "Upper and lower probabilities induced by a multivalued mapping", *Annals of Mathematical Statistics* **38** 325–339.

Fox, J, 1986. "Three arguments for extending the framework of probability", in: *Uncertainty in Artificial Intelligence*, L N Kanal and J F Lemmer (Eds.), Amsterdam: Elsevier Science Publishers B.V.

Hinde, C J, 1986. "Fuzzy Prolog", *J. of Man-Machine Studies* **24** 569–595.

Kanai, N and Ishizuka, M, 1986. "Prolog-ELF incorporating fuzzy logic", *Translations of Information Processing Society of Japan* **27** 4.

Lee, R C T, 1978. "Fuzzy logic and the resolution principle", *J. ACM* **19** 109–119.

Goodman, I R and Nguyen, H T, 1989. *Uncertainty Models for Knowledge-Based Systems*, Amsterdam: North-Holland Publishing Co.

Mukaidono, M, Shen, Z and Ding, L, 1987. "Fuzzy Prolog", *Congress Reprints of the 2nd IFSA Congress* **7** 452–455.

Nilsson, N, 1984. *Probabilistic Logic*, SRI Tech. Note 321, Menlo Park, CA.

Orlov, A I, 1980. *Problems of Optimization and Fuzzy Variables*, Moscow: Znanie.

Prade, H and Negoita, C V, 1986. *Fuzzy Logic in Knowledge Engineering*, Rhineland: Verlag TUV.

Preprints of the Second Congress of the International Fuzzy Systems Association, Tokyo, Japan, 1987.

Ruspini, E H, 1987. "Semantic modeling of imprecise and uncertain knowledge", *The Analysis of Fuzzy Information* **II**, J Bezdek (Ed.), Boca Raton, FL.

Shafer, G, 1976. *A Mathematical Theory of Evidence*, Princeton: Princeton University Press.

Spiegelhalter, D, 1985. "A statistical view of uncertainty in expert systems", in: *Proc. of Workshop on AI and Statistics*, W Gale and D Pregibon (Eds.), New York: Addison Wesley.

Sugeno, M, 1985. *Industrial Applications of Fuzzy Control*, The Netherlands: Elsevier Science Publishers B.V.

Togai, M and Watanabe, H, 1973. "Expert systems on a chip: An engine for real-time approximate reasoning", *IEEE Trans. on Systems, Man and Cybernetics* **SMC-3** 28–44.

Van Emden, M H, 1986. "Quantitiative deduction and its fixpoint theory", *J. of Logic Programming* **1** 37–53.

Watanabe, H and Dettloff, W, 1987. "Fuzzy logic inference processor for real time control: A second generation full custom design", *Proc. 21st Asilomar Conference on Signals, Systems and Computers*, Asilomar, CA.

Yager, R R, 1985. "Reasoning with uncertainty for expert systems", *Proc. 9th Int. Joint Conference on Artificial Intelligence*, Los Angeles, 1295–1297.

Zadeh, L A, 1979. "Fuzzy sets and information granularity", in: *Advances in Fuzzy Set Theory and Applications*, M Gupta, R Ragade and R Yager (Eds.), Amsterdam: North-Holland Publishing Co., 3–18.

Zadeh, L A, 1981. "Possibility theory and soft data analysis", in: *Mathematical Frontiers of the Social and Policy Sciences*, L Cobb and R M Thrall (Eds.), Boulder, CO: Westview Press, 69–129.

Zadeh, L A, 1983. "The role of fuzzy logic in the management of uncertainty in expert systems", *Fuzzy Sets and Systems* **11** 199–227.
Zadeh, L A, 1986. "Is probability theory sufficient for dealing with uncertainty in AI: A negative view", in: *Uncertainty in Artificial Intelligence*, L N Kanal and J F Lemmer (Eds.), Amsterdam: North-Holland Publishing Co., 103–116.

ECLECTICISM IN THE PARISH OF UNCERTAINTY

John Fox
Imperial Cancer Research Fund Laboratories, 44 Lincoln's Inn Fields, London WC2A 3PX

Along with the fashion for AI and expert systems has come an apparently religious argument about the nature of uncertainty. Broadly speaking the argument is between ecclesiastical and secular schools, which we may call the probabilistic and the eclectic schools. This argument has at times become strident. The worldly eclectics have been heard to make statements like "I know of no expert system project that either succeeded or failed on the basis of how it handled uncertainty" (R Davis, *personal communication*, 1985) while Mamdani and Efstathiou say that "... the current performance of expert systems is being limited by their capacity to cope with uncertainty" (Mamdani and Efstathiou, 1985, p. 283). The probabilist brethren, on the other hand, charge that the Eclectics are dangerous heretics proposing the use of *ad hoc* methods and/or reinventing wheels, thereby throwing away 300 years of thinking in the theory of probability.

Such quarrels are always damaging to the participants' causes. This one seems to be particularly so because it has distracted thoughtful people from deep problems which deserve study in both schools. It has long seemed to me, for example, that the preoccupation with quantitative details of the representation, revision and propagation of uncertainty has distracted us from understanding the limits of traditional mathematical theories of decision making and inhibited the development of new, more versatile approaches (Fox, 1983, 1988). Alessandro Saffiotti's paper develops the idea that a narrow view of uncertainty is restrictive, even in classical terms, that an eclectic approach is to be preferred, and that AI presents an important opportunity to break with restrictive classicism.

Now battles of interpretation have raged for at least three centuries within the probabilists' own ranks. The soldiers of the Bayesian church are a particularly quarrelsome lot, with a well-known reputation for ruthlessness in pressing their Belief. Their uncompromising message, that probability is the only possible representation for uncertainty (Cheeseman, 1985), has however cut through all doubt and uncertainty, offering a secure basis of ritual with which to solve all problems.

Saffiotti is saying, and I completely support him, that notwithstanding claims that there is only One True Way of conceiving uncertainty, a consideration of the history of physics, philosophy, politics – and AI – will reveal that there is rarely one true representation of anything. Different representations are useful for solving different types of problem. Numbers, like s-expressions, graphs, rules etc. were invented as representational tools, they are not revelation. Like all representations they are good for some things and less good for others. Numerical expressions are particularly good for representing magnitude, along with numerical operations for comparing and combining magnitudes. Non-numeric structures are good for representing complex objects, and for organizing, annotating and transforming those structures according to logics.

The argument for an eclectic view of decision making is not an argument for unprincipled pragmatism. I think that Saffiotti is as concerned with soundness and rigour as anyone, its just that he wants to extend that rigour to meta-level reasoning about uncertainty. To date the quarrel between members of different Ways has been grounded in intuitions – intuitions which are articulated in natural language, an uncertain tool. The symbolic, computational analysis which is made possible by AI languages may offer a less ambiguous method for clarifying the nature of uncertainty than the contemplation of our intuitions.

The clear corollary to this argument is that we need to be able to describe the various representations formally and explicitly, together with their properties and assumptions, so that we can reason about them and the contexts in which they are to be used (Fox, 1986). The existence of meta-level methods for reasoning about the nature of the uncertainty facing us, and what methods to apply in which kinds of situation, are fascinating possibilities. If they can be successfully developed I think they must significantly influence our conception of mathematical probability itself.

Extending this last point, the argument between eclectics and probabilists may be only one case of a general problem in symbolic reasoning; when should we (or a machine) use quantitative and when qualitative representations in problem solving? Various communities in AI and knowledge engineering are investigating qualitative reasoning. Some are interested in spatial reasoning (geometric representations "versus" topological representations of three dimensional structures); temporal reasoning (clocks and calendars "versus" temporal logics); and causal reasoning (physics based theories "versus" intuitive theories of causality) and so forth. The quantitative/qualitative divide is becoming particularly well articulated in the area of uncertainty – if we can resolve our differences we might shed some light outside our own parish.

Finally, since I have had the opportunity of seeing all the other contributions to this discussion I will allow myself one comment on them. It seems to me that to a greater or lesser extent most of the contributors to this discussion have moved towards what David Spiegelhalter calls in his piece an *ecumenical* spirit. This represents a major change of mood from a few years ago. I think that Saffiotti's paper, and the discussion which it has attracted, are a significant contribution to that change of mood. I congratulate him and say long may it continue.

References

Cheesemen, P, 1985. "In defence of probability", *Proceedings of the International Joint Conference on Artificial Intelligence*, Los Angeles, 1002–1009.

Fox, J, 1983. "Formal and knowledge based methods in decision technology", *Proc. 9th conference on Subjective Probability, Utility and Decision Making*, Groningen, The Netherlands, August 1983. Extended version in *Acta Psychologica*, 1984.

Fox, J, 1986. "Three arguments for extending the framework of probability", in: L N Kanal and J F Lemmer (Eds.), *Uncertainty in Artificial Intelligence*, Amsterdam: North Holland.

Fox, J, O'Neil, M, Glowinski, A J and Clark, D A, 1988. "Decision making as a logical process", *Proceedings of Expert Systems '88 Brighton.* Cambridge University Press 1988.

Mamdani, E H and Efstathiou, H J, 1985. "Higher-order logics for handling uncertainty in expert systems", *Int. J. Man-Machine Studies* **22** 283–293.