An AI view of the treatment of uncertainty

ALESSANDRO SAFFIOTTI

Dipartimento di Scienze dell'Informazione, Università di Pisa, c.so Italia 40, 56100 PISA, ITALY

Abstract

This paper reviews many of the very varied concepts of uncertainty used in AI. Because of their great popularity and generality "parallel certainty inference" techniques, so-called, are prominently in the foreground. We illustrate and comment in detail on three of these techniques; Bayes' theory (section 2); Dempster-Shafer theory (section 3); Cohen's model of endorsements (section 4), and give an account of the debate that has arisen around each of them. Techniques of a different kind (such as Zadeh's fuzzy-sets, fuzzy-logic theory, and the use of non-standard logics and methods that manage uncertainty without explicitly dealing with it) may be seen in the background (section 5).

The discussion of technicalities is accompanied by a historical and philosophical excursion on the nature and the use of uncertainty (section 1), and by a brief discussion of the problem of choosing an adequate AI approach to the treatment of uncertainty (section 6). The aim of the paper is to highlight the complex nature of uncertainty and to argue for an open-minded attitude towards its representation and use. In this spirit the pros and cons of uncertainty treatment techniques are presented in order to reflect the various uncertainty types. A guide to the literature in the field, and an extensive bibliography are appended.

Introduction

In the real world all knowledge is accompanied by a certain amount of uncertainty. Now perhaps more than ever scientific knowledge is pervaded by doubts, non-decisiveness, observations and laws which are probabilistic rather than absolute, and by the "artistic" component of reasoning. So called "common-sense knowledge" is even more deeply characterized by uncertainty, both in the raw information it depends upon and in the inferences that can be drawn.[1] Yet human beings are able to use this uncertain knowledge effectively to shape their model of reality, to take decisions, and to act. How this is done has been an intriguing concern of epistemological enquiry for centuries, as has the very nature of uncertainty itself. The ensuing debate has found new dimensions as well as new stimulus in the field of Artificial Intelligence (AI).

This paper is meant to be a fairly wide—though still not exhaustive—survey of AI approaches to the treatment of uncertainty. The window we open on these issues lets us appreciate a particular view with the so called "parallel uncertainty inference" techniques (defined in 1.3) in the foreground, both because of their popularity and because of their importance in the discussion;[2] in minor detail, we can see a vast forest of uncertainty types.

1 The need to represent uncertainty

The first steps in setting AI to cope with uncertainty lie in understanding what uncertainty really is, and in establishing techniques to formalize and use it in cognitive processes. The development of scientific research regarding these issues is discussed in this section, emphasizing the variety of meanings within the concept of uncertainty.

[1] Notice that uncertain information may come from inaccurate measures as well as from noise in the transmission processes, or from the intrinsic vagueness of the concepts used.

[2] This paper is a shortened English version of chapter 2 of (Saffiotti, 1988), which shows the results obtained by experimenting with the idea of merging uncertain and non-monotonic reasoning.

1.1 The past: from gambling-parlours to scientific pass-partout

The yearning for formalization, and the major development in mathematics in the XVII and XVIII centuries, did not stop mathematicians applying mathematics to mundane affairs. It was in this spirit that Pascal and Fermat conceived the probability calculus in the XVII century: it was just a way to study gambles. During the following century, when Giacomo Bernoulli and Thomas Bayes assembled the greater part of the probability calculus machine, probability calculus was still a "conjectare" (a conjecturing) to be applied to insurance company death-rate tables, while mathematics was a "scire" (a knowing), applied to physics.

Probability calculus, and the notion of probability itself, were really allowed to enter the mathematical-scientific forum only during the XIX century: it was then that Laplace described the development of science as a flow from ignorance to omniscience, a building up of theories which approach reality ever more closely. Probability calculus, which takes into account both partial knowledge and partial ignorance, then became the key concept of scientific development: it is the sole tool we have to formalize our progressive approach to reality.

Once accepted mathematically, probability calculus needed a precise definition of its foundations, in particular (and still controversially) the concept of probability itself (or, more generally, the concept of uncertainty).

Laplace defined the probability of an event e to be the ratio between the number of cases which are favourable to e (say m) and the number of possible cases (n):

$$p(e) = m/n$$

A more general definition is the "frequency" interpretation, due to Locke and Peirce: p(e) is now the limit of the relative frequencies of e while the number of tests performed tends to infinity. The test must be repeatable. This global, synthetical definition tends to dim the meaning of a single event, while it is well suited for those problems that are statistical in nature.

The meaning of probability for a single event plays a central role in the so called subjective interpretation: p(e) is the confidence a subject places—for reasons not to be dealt with—in the occurrence of e, and the conditions p is to satisfy are restricted to the mere "coherence principle" (the probability measures of a set of exhaustive alternatives must sum "sum to one" one). This interpretation, pioneered by de Finetti, and the ensuing debate, have found a new resonance in AI research: an outline of this debate will be given in section 2.5 and 3.1.

More recently the "logic" interpretation of probability has appeared: Rudolf Carnap (and also J.M. Keynes and H. Jeffrey) has focussed on the probability that, having observed that h is the case, an event e occurs, and called it c(e, h) ("confirmation of e by h"). What is being measured is not a statistical correlation, but rather the strength of a cause-effect law. This measure actually belongs to the scientific metalanguage. Carnap based a theory of inductive methods on it. A problem encountered when dealing with this interpretation deserves mention: "confirmation" is intensional in nature, that is it must take into account not just the extensions (the "meaning", after Frege) of e and h, but rather their intentions, the "sense" of e and h. In other words, e′ may be logically (extensionally) equivalent to e (i.e. the extensions—truth values—of c and e′ be the same in all interpretations), yet $C(e', h) \neq C(e, h)$, because the intensions of the two laws (the concepts they express) are not the same.[3]

1.2 The present: the thousand and one faces of uncertainty

The above review highlights the complexity hidden in the concept of probability. The history of this concept is not to be seen as a race towards a hypothetical "true" definition of probability, but as a step by step discovery of its various facets.

[3] A well known example of this problem is the "raven paradox": we cannot accept the observation of a red table to confirm the law "$\forall x \cdot raven(x) \rightarrow black(x)$" ("all ravens are black"), although it does confirm the logically equivalent law "$\forall x \cdot not(black(x)) \rightarrow not(raven(x))$" ("all the non-black things are not ravens"). We will also deal with the difference between intensionality and extensionality in section 5.1.

However this variety of uncertainty ideas did not prevent A.N. Kolmogarov from discovering a uniform axiomatic theory of probability calculus, in 1933, of which the above interpretations are models. Briefly his theory is defined over the power-set of the set $\theta = \{\theta_1, \ldots, \theta_n\}$ of elementary events (the "sample space"), and that the probability function

$$p: \mathrm{Bool}(\theta) \rightarrow [0, 1]$$

must satisfy, for each A, $B \subseteq \theta$

$$\begin{aligned} &\text{i. } p(A) \geqslant 0 \\ &\text{ii. } p(\theta) = 1 \\ &\text{iii. } p(A \cup B) = p(A) + p(B) \quad \text{if } A \cap B = \emptyset \end{aligned} \tag{1}$$

It is worth noticing that axiom iii (the "total probability principle") implies

$$p(A) + p(\sim A) = 1$$

The axioms above fit the classical interpretation as well as the frequency and the subjective interpretations (axiom iii is equivalent to the coherence principle). Things are not so plain for the logic interpretation: C(A, h) does not obey the laws $p(A \mid h)$ does. In particular, if we were to accept that

$$C(A, h) + C(\sim A, h) = 1$$

we should accept that an observation confirming A to a degree of, say, x also confirms the negation of A to some degree (namely $1 - x$). The wheel turns the way the raven paradox does. The solution is[4] to keep confirmation distinct from disconfirmation.

Returning to AI we must keep in mind that the point is not so much finding the best axiomatization or the best interpretation for uncertainty, as the pragmatic goal of adequately representing uncertainty. We imagine, in the spirit of AI, that the aim will be to identify the many facets of uncertainty and to represent them so that the knowledge they embody can be used effectively.

1.3 The future: AI research on uncertainty representation

We mentioned before that any system which is to face real world problems, either natural or artificial, must confront some form of uncertainty. The first way to do this may merely be to ignore the uncertainty. One takes a model of reality in which particular assumptions make all knowledge certain; that an uncertain measure is exact; that a missing datum has a "default" value; that a probabilistic rule is universally valid; that a non-demonstrably true proposition is false, and so on.

A drawback of this technique is the need to revise our knowledge if the assumptions prove to be false. This issue is central is non-monotonic-reasoning, and we won't pursue it here.[5] But even if assumptions need not be revised, there is the risk of formalizing a problem which proves to be so different from the original one as to be quite useless, or of losing valuable information from discarding uncertain knowledge. This often forces AI to explicitly recognize and handle the uncertainty that is present. Still, there are cases in which the most effective solution—or perhaps the only one—is just to ignore uncertainty; actually, in every real-world problem some aspects of uncertainty are inevitably ignored. It is interesting to observe that there are, symmetrically, problems of a deterministic nature in which uncertainty is introduced as a means of coping with complexity (e.g. chess).

When uncertainty is to be dealt with, however, the possible solutions are as many and various as its interpretations are. Many of these solutions share the attitude of viewing the knowledge and the uncertainty about it as two independent entities, and so treating them by means of two distinct loosely-coupled processes: the reasoning process handles knowledge as if it were exact, while a

[4] In 2.3, we'll see how AI has coped with this problem.
[5] The interested reader is referred to the special issue of *Artificial Intelligence* (vol. 13, 1980).

"parallel uncertainty inference"[6] process accompanies it, computing the uncertainty affecting each newly derived fact. This uncertainty is in its turn usually based on the uncertainty affecting the facts used to derive the new fact.[7] The way uncertainty is represented and processed by the parallel uncertainty inference process is a distinguishing characteristic of the different techniques.

There are so many different uncertainty management techniques that it would be convenient to have a common framework in which they can be formalized (we shall give further motivation for it in 6); Thompson (1985) proposes a clean general paradigm in which to formalize them. This is structured into four parts:

a. the base elements on which the theory is defined: they will typically consist of an algebra of statements describing the object domain, and some certainty and utility functions defined over this algebra;
b. the observation reports, that is how the new evidence is represented in the theory;
c. the updating mechanism, defining how the transitions from one state of certainty to another are performed;
d. the decision mechanism, which reaches a decision based (also) on the state of certainty.

The updating mechanism (point c) is pivotal to any theory. It implements what Domotor (1985) called the "probability kinematics",[8] that is, the movements of the probability masses among the statements of our algebra.

Some of the most popular evidential reasoning theories in AI will be presented in subsequent sections. Thompson's paradigm will be helpful in the presentation.

2 Bayesian methods

By "Bayesian methods" we mean those techniques based on the classical probability theory and the use of Bayes' theorem. The following sections deal with the definition, the meaning and the use in AI of this theorem. The debate raised about them—in AI as well as within the scientific-philosophic community—is briefly presented in section 2.5.

2.1 Bayes' theorem

Bayes' theorem dwells on how probabilities attached to a set of (exhaustive and mutually exclusive) hypotheses $\theta = \{\theta_1, \ldots, \theta_n\}$ are to be revised in light of a new evidence e. In its most popular form, it states that

$$p(\theta_i \mid e \,\&\, H) = \frac{p(\theta_i \mid H) * p(e \mid \theta_i \,\&\, H)}{p(e \mid H)} \tag{2}$$

where all the probabilities are conditioned by some former information H. H could consist of the mere hypotheses of the problem, but usually it consists of the evidence $\{e_i, \ldots, e_m\}$ currently obtained.

The importance of this theorem lies in its expressing a measurable proportionality (the likelihood $p(e \mid \theta_i \,\&\, H)$) between the probability of θ_i before and after the acquisition of the new evidence e. From an epistemological point of view this means formalizing an inductive behaviour, which Lindley (1971) asserts models the greater part of adult learning.[9] For our purposes it is sufficient to

[6] We are indebted to Cohen (1985) for this definition.

[7] Famous expert systems like MYCIN (Shortliffe-Buchanan, 1975), PROSPECTOR (Duda et al., 1976), or expert systems shells like EMYCIN (vanMelle, 1980) or EXPERT (Weiss-Kulikowski, 1979), can be fitted into this paradigm. A wit example of this philosophy can be found also in the technique proposed by E. Shapiro (1983) to enrich logic programming with an uncertainty handling mechanism.

[8] He defines a clever general theory in which the most popular probability kinematics (Bayes', Gibbs', Jeffrey's, etc.) can be formalized.

[9] He also emphasizes the epistemological danger of absolutely excluding a possibility for something to be true (or false): every learning is inhibited in (2) if $P(\theta_i \mid H) = 0$.

notice that Bayes' theorem is a good candidate for the updating mechanism of an uncertain reasoning system.

2.2 *The use of Bayes' theorem in AI*

First of all, it is worthwhile to notice that very few of the AI "Bayes-based" systems do rigorously apply Bayes' theorem: most of them, as we are about to see, actually make use of approximate formulas derived from it. With this in mind, we can describe the Bayes-based approach to uncertain reasoning in AI following Thompson's paradigm: it will usually consist in defining

a. the world model base elements:
   1. a database of statements $\theta_1 \ldots \theta_n$
   2. a probability function p defined over this database,
   3. a utility function u with a statement and possible action as arguments
b. the evidence acquisition mechanism; usually limited to assigning probability values to statements;
c. the updating mechanism; usually based on Bayes' theorem or on some apprximation of it.
d. the decision mechanism, which usually chooses the action A_j maximizing the expected utility

$$U(A_j) = \sum_i u(A_j, \theta_i)p(\theta_i)$$

Often, in AI applications, formulations of Bayes' theorem other than (2) are useful:

$$p(\theta_i \mid e_h \,\&\, e_k) = \frac{p(\theta_i) * p(e_h \mid \theta_i) * p(e_k \mid \theta_i)}{p(e_h) * p(e_k)} \tag{3}$$

is computationally expedient in the most restrictive hypothesis that the incrementally collected pieces of evidence $e_1, \ldots, e_m$ are both independent with respect to the entire evidence population and independent with respect to the conditioning event θ_i.

$$0(\theta_i \mid e) = \lambda * 0(\theta_i) \tag{4}$$

where

$$0(\theta_i) = \frac{p(\theta_i)}{p(\sim\theta_i)},\ 0(\theta_i \mid e) = \frac{p(\theta_i \mid e)}{p(\sim\theta_i \mid e)},\quad \text{and}\quad \lambda = \frac{p(e \mid \theta_i)}{p(e \mid \sim\theta_i)}$$

offers a solution to the main problem of the Bayesian approach; this problem, is that the method requires a bulk of data (a-priori and conditional probabilities) which is often hard—if not impossible—to come by. (4) requires only a small number of "likelihood ratios" λs; these numbers can be extracted directly from the expert's (subjective) opinions, at the cost of possible inconsistencies in the values supplied (hence escaping from the formal Bayesian apparatus). Techniques of this kind, called "subjective Bayesian", are explored by (Duda et al., 1976) in the context of the PROSPECTOR expert system, and by (Pednault et al., 1981).

2.3 *The MYCIN model*

Switching from rigid application of Bayes' theorem to its computational approximations, it is perhaps instructive (and de rigueur) to glance at the model that Buchanan and Shortliffe (1975) developed to compute the conditional probabilities $p(\theta_1 \mid e)$.

They first noticed that a conditional probability $x = p(\theta_i \mid e)$ can be seen as a weighted rule stating "if e is the case, then θ_i is, with certainty x". Then, if $E = e_1, \ldots, e_m$ is the evidence gathered, the AI problem is to compute $p(\theta_i \mid E)$, for each i, merely on the basis of the values $p(\theta_i \mid e_j)$, $j = 1, \ldots, m$; that is to compute the certainty of each conclusion θ_i drawn from arbitrary applications of weighted rules, on the basis of their weights.

Secondly, they noticed that this interpretation of conditional probability did fit Carnap's confirmation theory (see in 1.1), so sharing the need to keep $p(\theta_i \mid e)$ distinct from $p(\sim\theta_i \mid e)$. Then, they defined

$$MB(\theta_i \mid e)$$

to be the "measure of belief" in θ_i given the evidence e, while denoting with

$$MD(\theta_i \mid e)$$

the corresponding "measure of disbelief", and combined them into the so called "certainty factor".

$$CF(\theta_i \mid e) = MB(\theta_i \mid e) - MD(\theta_i \mid e)$$

These CF's don't obey the probability calculus axioms; for instance

$$CF(\theta_i \mid e) + CF(\sim\theta_i \mid e) = 0 \tag{5}$$

contravening axiom iii of (1), as is the case for confirmations.

Unfortunately, (5) also shows ways in which CF's can't overcome difficulties with probability. First, the certainties of the hypothesis and of its opposite are still linked, and to assert something about one forces us to assert something about the other (though in the negative). Also, a null CF may come either from $MB = MD = 0$ (total ignorance), or from $MB = MB \neq 0$ (contradiction), failing to distinguish the two cases. These problems, further discussed in 2.5, arise from the failure to represent ignorance; a possible answer will be seen in the next section.

Buchanan and Shortliffe applied these ideas to the construction of a rule-based system for medical consultation—the famous MYCIN. The θ's and the e's were taken to be respectively the diagnosis and the symptoms, while a whole machinery was devised alongisde the rules to compute the CF of each diagnosis given a set of symptoms. Buchanan and Shortliffe (1984) give an extensive account of MYCIN and of its developments.

The MYCIN model is an ad-hoc technique; yet, it is—at least in principle—applicable to every problem that turns out to be isomorphic to the MYCIN one, provided that the statistical independence assumption implicit in the model (Adams, 1976) is granted.[10] One must also be conscious of the weaknesses of the CF's: they rapidly converge to 1 as the number of supporting rules increases; the hypothesis with the highest CF is not guaranteed to be the most probable one; and the insensitiveness of MYCIN to a change in the CFs of its rules is puzzling (Buchanan-Shortliffe, 1984, ch. 10). The reasons for the excellent performance MYCIN shows in spite of these limitations are likely to be found in the small number of rules usually applied to reach a conclusion.

2.4 *Other Bayes-based methods*

After the MYCIN experience, a host of models and implementations more or less based on the Bayesian method appeared in AI.

In the expert system SPERIL, Ishizuka, Fu and Yao (1981) the MYCIN CF technique is merged with Zadeh's fuzzy sets (see 5.2), so managing to represent the vague knowledge they were dealing with, while preserving the simplicity and modularity which are characteristic of rule-based systems.[11]

Friedman (1981), noticing that the inferences performed by MYCIN are limited to the "confirmation" rule, extends the MYCIN model to allow "modus ponens", "denial" and "modus tollens" derivations.[12] In addition, he proposes to compute the rules' CFs (seen as weights labelling the search graph edges) from information acquired by the system, rather than retrieving them from

[10] This can often be accomplished by grouping all dependent fragments of evidence together.

[11] In subsequent works, they revert to the use of Dempster-Shafer's theory, again extended to deal with fuzzyness (Ishizuka, 1987).

[12] Remind that these four inference rules respectively declare that $B, A \rightarrow B \vdash A$, $A, A \rightarrow B \vdash B$, $\sim A, A \rightarrow B \vdash \sim B$, $\sim B, A \rightarrow B \vdash \sim A$.

experts' knowledge. An autonomous computation of the rules' CFs is also suggested by Cheeseman (1984), who proposes a statistical method based on Bayes' formula.

Finally, mention should be made of the "convex" Bayesian approach: it mimics the classical one, except for the probability functions which prove to be a convex set of functions. Bayes' theorem mow maps convex sets (of functions) into convex sets in response to new observations. This approach, however, does not seem to have found a great deal of application in AI.

2.5 *The hornet's nest*

Long before its use in AI, the Bayesian approach to uncertainty treatment had been the focus of wide debate. AI was destined to become a resonance chamber for this debate, and a source of new expectations.

A first batch of criticisms concerned the amount of data necessary for Bayes' theorem to work. This is a computational adequacy criticism, and it stimulated research into alternative computational models more than philosophical debate (but see Shafer, 1987). Epistemological adequacy criticisms, on the other hand attack[13] the use of a single value to represent uncertainty on different fronts. On the one hand, it seems to be too reductive to sum up all our knowledge about (and all the different aspects of) uncertainty in a single, linear measure; on the other, the single value representation does not seem to be able to represent our ignorance. All of the three statements

"The probability of A being the case is about 60%"
" „ „ is between 40% and 80%" (6)
" „ „ is exactly 60%"

are collapsed into the third one when represented as

$$p(A) = 0.6$$

The answers given by Bayes' advocates usually follow the method of "extending the conversation" (Lindley, 1971): any kind of uncertainty, they affirm, can be modelled in a Bayesian framework, provided that new variables were introduced in the conversation. So, if we set X to assume as values the possible probability values for A, we'll be able to represent the (6)'s as, respectively,

$$p(X(A) = x) = \begin{cases} 0.5 & \text{if } x = 0.6 \\ 0.25 & \text{if } (X = 0.5 \quad \text{OR} \quad x = 0.7) \\ 0 & \text{else} \end{cases}$$

$$p(X(A) = x) = \begin{cases} 0.2 & \text{if } 0.4 \leqslant x \leqslant 0.8 \\ 0 & \text{else} \end{cases}$$

$$p(X(A) = x) = \begin{cases} 1 & \text{if } x = 0.6 \\ 0 & \text{else} \end{cases}$$

Bayes' theorem, in the formulation

$$p(A \mid e) = \sum_{x \in U} x * \frac{p(X = x) * p(e \mid X = x)}{p(e)}$$

models our intuition according to which evidence is to modify vague prior probabilities (in which the X distribution function is out of focus) more than sure ones (in which the X distribution is more concentrated).

It is arguments like these that allow Bayesians to offer challenges like "anything that can be done by these [non-probabilistic] methods can be *better* done with probability" (Lindley, 1987, p. 20; our

[13] Another batch of criticisms applies more generally to the use of numbers in representing uncertainty, and we'll deal with it in 4.1.

emphasis). However, it is worth noticing that these answers actually shift the adequacy problem from the epistemological level to the computational one (both representational and computational complexities are heavily increased by extending the conversation), therefore sounding more harmonious to the philosopher's ears than to the AI designers'.

3 Dempster-Shafer theory

Dempster-Shafer theory is a recent attempt, built on solid mathematical ground, at formalizing an uncertainty calculus in such a way that it does not suffer from the adequacy problems seen in the previous section. The structure of this section mimics that of the one above, ending by comparing the Dempster-Shafer theory with the Bayesian one.

3.1 *The theory*

Dempster-Shafer theory came into being in the sixties from the work of Arthur Dempster (e.g. 1968), and was then put into a suitable form for finite domains by a student of his, Glen Shafer (1976). It is defined over the power-set of the set of base elements $\theta = \{\theta_1, \ldots, \theta_n\}$ called the "frame of discernment" (or, equivalently, over the set of all the disjunctions of a set of statements). A mass probability function m, satisfying

$$\mathrm{m}(\phi) = 0$$

$$\sum_{A \subseteq \theta} \mathrm{m}(A) = 1$$

assigns a value in [0, 1] to every subset of θ, so distributing a unit of "probability mass" among them. It is worth noticing that m being directly defined on Bool(θ) allows us to specify the available information at just the level of detail where it belongs, being non-committal on the levels we are ignorant about. In particular, m(θ) will be the probability mass we cannot commit to any smaller subset of θ: it will actually represent our ignorance.

Some derived measures are useful:

$$\mathrm{spt}(A) = \sum_{B \subseteq A} \mathrm{m}(B)$$

gives a measure of how much of the overall probability mass in some way supports A, and

$$\mathrm{pls}(A) = 1 - \mathrm{spt}(\sim A)$$

which measures the overall probability mass not supporting ~A. The probability of A is bounded within the interval [spt (A), pls (A)], the width of which measures how exactly we can specify p(A). In other words, this width measures our ignorance about A.

An example will help clarify the meaning of these measures.

The 'Francesco's image' problem:

The puzzling question is whether Francesco will dress in black (B), white (W) or polka-dot (P). The problem domain can be depicted as in Fig. 1: with edges standing for propositional implications.

Let's suppose the only available information is that Francesco's washing machine is out; then, he is unlikely to dress in white, while we still have no means of discriminating between B and P. A probability mass function could then assign the 80% of the probability mass to B ∨ P (i.e. to ~W), being non-committal about the remaining 20% (i.e. assigning it to θ). We'll get the situation shown in Fig. 2a. Note that we are completely ignorant about B and P separately (we only know that the probability lies in [0, 1]; which is not very useful), while having some knowledge about B ∨ P; W is not impossible, but it is fairly implausible.

Now, were we to hear Francesco dislikes dark things, we would be able to reject hypothesis B. To do this, we can imagine lowering the mass committed to θ upon W ∨ P, and that committed to B ∨ P upon P, so excluding B (see Fig. 2b). We'll see how the theory formalizes this kinematics in a moment.

It is worth noticing how the [spt, pls] intervals for P and B narrow in response to the new evidence. In 3.4, we'll again refer to this example, comparing Dempster-Shafer and Bayesian techniques.

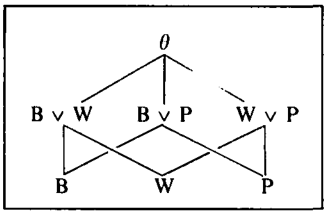

Figure 1

m	spt	pls	hypothesis	m	spt	pls
0	0	1	B	0	0	0
0	0	0.2	W	0	0	0.2
0	0	1	P	0.8	0.8	1
0	0	1	B ∨ W	0	0	0.2
0.8	0.8	1	B ∨ P	0	0.8	1
0	0	1	W ∨ P	0.2	1	1
0.2	1	1	θ	0	1	1

(a) (b)

Figure 2

Let us now look at how probability kinematics is modelled by Dempster-Shafer theory."Dempster's combination rule" combines two mass assignments m_1 and m_2 (usually representing a prior certainty state and a new evidence respectively) into a third one, m, by

$$m(A) = \frac{1}{K} \sum_{X \cap Y = A} m_1(X) * m_2(Y) \tag{7}$$

with

$$K = 1 - \sum_{X \cap Y = \phi} m_1(X) * m_2(Y)$$

K is a normalizing factor which measures how much of the mass would be committed to the empty set (an event that is bound to happen whenever m_1 and m_2 assign mass to two disjoint sets); K so acts as a measure of how much m_1 and m_2 are conflicting. If $K = 0$, m_1 and m_2 are said to be "completely contradictory". Notice that (7) is both commutative and associative, so the order in which information is collected is irrelevant.

A graphical representation of the combination rule is shown in Figure 3, where m_1 and m_2 are represented as partitions of the unitary segment, and m(B) is the total area of all the rectangles committed to B as in figure

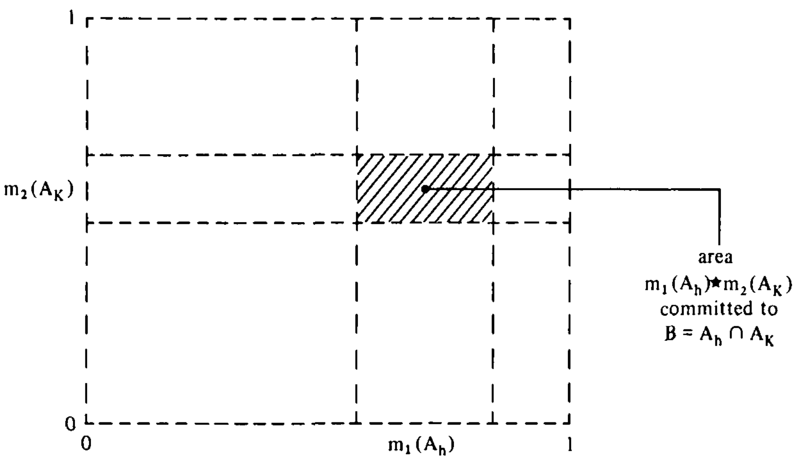

Figure 3

3.2 The use of Dempster-Shafer theory

According to Thompson's paradigm, a Dempster-Shafer based system must define:

a. the base elements, consisting of
   1. the algebra of statements θ,
   2. the probability mass functions m,
   3. an utility function u;
b. the evidence acquisition mechanism, i.e. the translation of evidence into a mass distribution over θ;
c. the updating mechanism, applying Dempster's rule;
d. the decision mechanism, based on the spt and pls measures and on the u function.

Some points must be emphasized. First, the definition of θ must be carefully tailored to the application domain, so that every new piece of evidence can be easily translated into a mass distribution over θ. The translation process may be a burden both from computational and design points of view. An interesting example of sensor supplied data translation can be found in (Garvey et al., 1981). On the other end of the reasoning process, the problem of how to design the decision mechanism is still an open problem: no general technique has yet been developed to cope with probabilities which are intervals rather than points (this is one of the thorns of the rose of representing ignorance).

Notice that theory defines only the probability kinematics due to evidence acquisition, leaving propagation through logical constraints to be defined by the implementation. (Dubois-Prade, 1985) discuss the problem, and (Garvey et al., 1981) propose a propagation mechanism based on natural deduction. This latter proposal is also an example of how Dempster's rule can be used to combine knowledge from multiple sources (data from several devices as well as opinions from different experts), even providing a measure (K) of the total agreement. (Ginsberg, 1981) points out how Dempster's rule can be effectively used in problems in which knowledge may have to be retracted. Finally, the computational complexity of Dempster's rule can be reduced to linear time under fairly weak hypotheses (Barnett, 1981).

3.3 Methods based on Dempster-Shafer theory

Though young, Dempster-Shafer theory has already excited the AI community. (Lu-Stephanou, 1981) used it to combine certainties in a kind of set-theoretic knowledge-based system, while solving the propagation problem by a plain multiplicative mechanism; it also proposed a certainty decomposition rule that is an inverse of Dempster's rule. (Strat, 1984) extended the theory to deal with continuous variables, adequately redefining the domain θ as well as the combination rule. (Ginsberg, 1981) applied the theory to semantic networks: the links are labelled by certainties, the Dempster's rule is the sole mechanism to both propagate and combine the certainties through the network. As noted above, Ginsberg addresses non-monotonic reasoning problems as well; a wider debate about the marriage between uncertain and non-monotonic reasoning techniques is to be found in (Saffiotti, 1988).

3.4 A comparison with the Bayesian techniques

First of all, we must notice that Dempster-Shafer theory includes the Bayesian one.[14] Shafer (1976) shows that the latter is a particular instance of the former when the probability masses are constrained within single points. To see how this constraint limits the Bayesian theory it would be useful to re-examine the "Francesco's image problem" in a Bayesian framework:

[14] (Garbolino, 1986) shows how both of them can be modelled in a common generalized Bayesian probability kinematics framework (see also Domotor, 1985).

hypothesis	p(·)	p(.le)
B	0.4	0
W	0.2	0.33
P	0.4	0.66
B ∨ W	0.6	0.33
B ∨ P	0.8	0.66
W ∨ P	0.6	1
θ	1	1

Figure 4

The original knowledge is represented by the probability function

$$p(W) = 0.2, \; p(B) = p(P) = 0.4$$

where, in spite of our ignorance, we have been obliged to assign exact probability values to B and P (arbitrarily assumed to be equally probable), and to convert the non-decisiveness or our argument against W into a small probability W itself. Applying Bayes' formula, the new evidence e being Francesco's comment about dark clothes, we get the results shown in Figure 4; notice that both the prior and the conditional probabilities of W are just half the ones of P's, while none of the available informations directly supported W. In the Dempster-Shafer analysis, the probability of W was bounded within the [0, 0.2] interval, and this interval had not been altered by the evidence about B.

To summarize, we can observe that

(1) Dempster-Shafer theory allows one to specify the available knowledge at just the right level of detail, while this knowledge—whether or not it allows us to—has to be stated at the ground level in the Bayesian approach;
(2) through the [spt, pls] interval width, Dempster-Shafer's representation is able to directly model that "certainty about certainty" which the Bayesian theory was criticized for not accommodating. However, this does not alter the "stiffness" of the certainties when they are confronted with new evidence; the vagueness of a probability value and its sensitivity (its disposition to learn from experience) turn out to be two distinct concepts;
(3) in Dempster-Shafer theory there is no rigid constraint between the probability mass committed to A and that committed to $\sim$A;
(4) in both Dempster-Shafer and in Bayes' theory small probabilities behave in a qualitatively different way from null ones (Dubois-Prade, 1985).

While the comparison above may suggest the superiority of Dempster-Shafer theory (at least for AI) we emphasize with Garbolino (1986) that Dempster-Shafer's and Bayes' theories are not to be seen as rivals, but as alternatives, both belonging to a family of probability kinematics theories among which we are allowed to choose the one that best fits our problem. Bayesian methods, for instance, might be the most convenient choice, from a computational point of view. Shafer himself, in comparing Bayes' and Dempster-Shafer's theories, defends a very balanced position in which the kind of evidence which is available is considered in the evaluation of the theories (Shafer, 1987). In section 6 we shall also insist on a permissive position.

4 Explicit representation of uncertainty

The techniques discussed in the previous sections—and actually most of the existing ones—make use of numbers to represent strength of belief, and of mathematics to manipulate the numbers. A new idea is to consider uncertainty as a kind of knowledge (about knowledge), and to use AI oriented techniques to represent and process it (Fox, 1986). Cohen's endorsement theory seems to be the most developed one to date; it will be presented here, alongside the debate about the adequacy of numerical methods and explicit uncertainty representation which inspired it.

4.1 The critique of the numerical method

It was during an unduly wet day that Mr. Pouring and his meteorologist friend Mr. Rain resolved to build a weather forecasting expert system. The problem quickly arose to attach certainty factors to the rules, and Mr. Rain turned into a circumspect mood.

"Whatcha mean, man?—he asked—That rule is only statistically true, but even about statistics there are lots of different opinions. Anyways, it's the best assumption we can make, seeing we don't have other information; provided this—he pointed to a scribble on the sheet—is not taken too seriously. Besides, ... uhm ... well, forget it: you wouldn't catch it. Anyway, all things considered, I, for myself, well, I guess I'd give it 6 to 10. Maybe."

"Excuse me, sir—Mr. Pouring added—but this rule deals with sudden frosts. So I suggest it would be better to boost its strength a little bit, just to be sure that the system warns us of the danger. What about 8?"

"God knows ... Why, let's do it your way!"

Mr. Rain is not the only expert feeling uncomfortable with numbers, when required to sum up all his knowledge about the truth of a statement (its pros and cons, its applicability conditions, his personal inclinations, and so on) into a number. Besides, the numbers so built will often not fit coherently into probability theory or other normative ones. The problem, first glimpsed in (McCarthy-Hayes, 1969) but clearly pointed out only at the 6th IJCAI (Davis, 1979), lies in these numbers combining a wide and structured body of knowledge, which remains inaccessible to the system which is to use it, as it is still implicitly handled by its arithmetic algorithms. On the other hand, were this knowledge explicitly represented and used as meta-level knowledge, it would provide us with valuable information about the use of the object-level knowledge we have (or we miss).

An advantage proclaimed for the numerical representation is the relative ease of comparing and ranking hypotheses. Number critics, however, say these comparisons are deceptive: one can assign probability estimates to wet feet being a cause of cold, and to the soul being reincarnated, but a comparison between them is surely misleading. Just as it makes little sense to compare a hypothesis supported by a weak argument with one which is supported by both strong evidence against and still stronger evidence for: we feel that the conflict on the second hypothesis has to be resolved more "intelligently" than by an arithmetic simplification. Moreover, the precision of numbers is usually high compared with the precision our knowledge justifies (yet we feel sceptical about a system that is too insensitive to these numbers: cf. Buchanan-Shortliffe, 1984, ch. 10). Besides, it is not clear when a difference in the certainty values is too little to be meaningful. Finally, when comparing hypotheses, the system must keep in mind its reasons for doing it.

Other arguments against the use of numbers to represent uncertainty rest on the fact that human beings seem not to use them when coping with uncertainty; Fox (1987) observes that we make use of qualitative terms to express logically distinct uncertainty states ("possible", "plausible", "probable"); it seems to be hard to compress these states into a flat numerical scale; Tversky and Kahneman (1974), analyzing a number of results in experimental psychology, suggest that humans use a small number of heuristics to reduce uncertainty, and notice that they fail to obey the laws of probability.

In spite of these faults, however, when the semantics of the uncertainty in the problem at hand is clear and well fitted to a numerical representation, the use of numbers may turn out to be the most effective solution. The criticisms imply we must investigate the very nature of uncertainty in a deeper way, and develop techniques to cope with those cases in which a plain, quantitative, synthetic treatment of uncertainty does not suffice.

4.2 The meta-knowledge interpretation of uncertainty

One of the most innovatory characteristics of AI is its concern with representing and using knowledge in as explicit a form as possible. This principle does not seem to have been applied to uncertainty, for which the implicit numerical methods discussed above have usually been the only possibility.

The explicit uncertainty representation idea consists of: (1) representing uncertainty as knowledge about knowledge (about its validity, its limitations, its pros and cons, etc.), and (2) working on this meta-knowledge making use of AI knowledge representation and processing techniques.

While the meaning of representing uncertainty as meta-knowledge is easily caught (we can imagine—not so far from Cohen's idea—appending to each piece of knowledge a frame containing all the information about its certainty), the second point deserves a little explanation.

My friend Giancarlo is a school-books salesman. My faith in his honesty is, though high, not complete (let's say 80%); in fact, although he has a generally honest nature, he lies about his school-books. Once I must employ my knowledge about Giancarlo's honesty—e.g. to measure the plausibility of something he has said—it would be too reductive if I was to limit myself to consider every utterance of his as equally probable (80%); instead, I must keep in mind the reasons which I base this 80% probability on, and consider what the topic of our conversation is all about (viz. if it deals with school-books or not). All this knowledge must be represented and used in our reasoning.[15] Moreover, an explicit representation of uncertainty may give valuable information on how the uncertainty itself can be reduced (by analyzing its sources) or on when it is possible, for the problem at hand, to ignore it. Uncertainty plays an active role, opposite to its passive one when seen as a numerical measure.

Finally, the explicit representation allows advanced explanation facilities, distinguishing for instance uncertainty of a conclusion due to weak evidence, from uncertainty originating in a too general rule. This kind of information is of great value in the difficult phase of knowledge revision.

It must be clear that the explicit representation of uncertainty is not so much an alternative to the numerical ones, as a complement: numerical or symbolic representations will be used (eventually together) depending on the kind of uncertainty we are talking about. This permissive position will be further expanded in section 6.a.

4.3 Cohen's model of endorsements

Originating in an interest in human reasoning about uncertainty, Cohen's proposal about uncertainty management in AI (published as Cohen, 1985) seems the most promising non-numerical approach to the problem to date. He defines an "endorsement" to be a structured object attached to a statement in which the reasons for believing or disbelieving in that statement are explicitly represented. In a telling analogy, Cohen describes the endorsements as the stamps and notes by which bureaucrats (rules) endorse a paper (a piece of knowledge): all that is needed to determine the validity of the statement can be represented in an endorsement, and inspected, modified or added by a rule concerned with it. The semantics of these endorsements is operationally defined by stating the way they are combined, propagated and ranked: this semantics, as well as the endorsement set itself, will depend on the application domain, a distinctive characteristic of Cohen's model compared with the ones we have seen above.

Cohen's model is implemented as a program (SOLOMON). The implementation of the model is critical, because 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.[16] Endorsements are attached in SOLOMON to every kind of knowledge (to inference rules, to program tasks, to data, to conclusions, etc.); some of these are of general use, others are problem specific, user-defined endorsements.

As an illustration, let's consider the rule Mr. Rain and Mr. Pouring were talking about a few lines above: the endorsements in Fig. 5 could be chosen to represent Mr. Rain's arguments about the rule's certainty;[17] then, of course, the way of using these endorsements should be defined in the expert system. It is this way of using endorsements, viz. their propagation, combination and ranking,

[15] Lindley (1987) emphasizes that Bayesian theory actually takes the whole context into account: the conditioning H in (2) above can in general embody all the information relevant to a problem. While this is a most elegant philosophical argument, it seems hardly acceptable from a practical AI point of view: where is this flood of p(e|H)'s supposed to gush from?

[16] The way endorsements are used to affect behaviour is indeed part of their operational semantics.

[17] Were we to use an hybrid (i.e. both symbolical and numerical) uncertainty representation, we could specify a relative frequency for e1, for instance by

e1. (correlational 70% vague)

e1	correlational	[the premise and the conclusion are statistically correlated, without a declared cause-effect relationship; cf. frequence vs. logic interpretation of probability in section 1.1]
e2	default	[best assumption in the absence of further date]
e3	(supportive P1)	[the scribble (P1) is not a necessary premise]
e4	dangerous	[the conclusion deserves particular attention]
e5	etc.	

Figure 5

that is the most obscure aspect of Cohen's proposal: there is indeed little understanding, once we have explicitly represented the uncertainty our knowledge is affected by, of how this is to affect conclusions we draw. SOLOMON actually propagates and combines endorsements by simply accumulating them, but Cohen himself insists on the need to devise more sophisticated techniques. On the other hand, how these endorsements can be ranked is not clearly understood either, perhaps reflecting the difficulty found by humans when balancing alternative hypotheses (difficulties often admittedly hidden by the numerical methods). These problems with the endorsement model, are to be seen as a sign of our poor understanding of uncertainty as a source of knowledge, rather than mere drawbacks, and do not compromise the validity of Cohen's intuition.

4.4 *The progress of Cohen's model*

The most problematic point of the model of endorsements, we said, is the way they propagate and combine (i.e. their kinematics), and the way they can be ranked. These are just the directions which Cohen has been investigating during the last few years.

The combination problem has been investigated in the framework of a plan-recognition program HMMM (Cohen-Sullivan, 1985). It elaborates on the idea of developing "semantic combination rules" for endorsements (Cohen, 1984): these rules attach (or delete) endorsements to hypotheses based on the presence or absence of other endorsements, taking into account their intended semantics.[18] Cohen and Sullivan also notice that numerical degrees of association between an endorsement and the knowledge it endorses may be helpful, moving away from Cohen's consciously drastic position of using only qualitative, symbolic reasons for believing or disbelieving.

As for the ranking of endorsements, no well-developed idea yet seems to have sprung up, but a promising research direction may be to partition the hypotheses into implicitly ranked equivalence classes on the basis of their endorsements, through flexible and possibly context-dependent partitioning functions.

5 Other uncertainty representation techniques

Many other approaches to the treatment of uncertainty, which do not fit the "parallel uncertainty inference" schemata (the main focus in this paper), have been developed. Among them, three seem to be more general and theoretically interesting: the use of non-standard logics, the "fuzzy" techniques, and the methods in which uncertainty is not explicitly handled, while the system's behaviour is nonetheless influenced. They will be briefly presented in this section, while references to other techniques will be given in the bibliographical notes.

5.1 *Non-standard logics*

Though a target of a wide debate, logic (viz. first order predicate calculus, PC) is undoubtedly a widespread basis for knowledge representation in AI. Developed for formalizing well structured

[18] Mind that endorsements are just tokens, which assume their meaning by the way they are used in the system (not unlike numbers, indeed).

fields like mathematics, it often showed its weaknesses when confronted with real world problems; these weaknesses have usually been overcome by: (1) somehow extending it (still maintaining the validity of PC's theorems), or (2) modifying it (obtaining systems in which some of the tautology of PC is no more valid). As examples two non-standard logics that lend themselves to formalizing uncertain knowledge will be illustrated.

The multi-valued logics

These logics (MVL's) maintain the syntax of PC, but alter the semantic interpretation, allowing an arbitrary number (n) of truth values, usually standardized over the [0, 1] interval, instead of the usual two values (true, false). MVL's are an example of the second approach to altering logic: the "tertium non datur" law is no longer a tautology when $n > 2$.

MVL's were first devised as three-valued logics in 1920 by J. Lukasiewicz;[19] the same author extended them to n truth values ten years later, by defining a "multivalued model" to be a structure $M = \langle D, F \rangle$ in which the mapping F maps each predicate symbol not just to a relation on the semantic domain D (as it is the case for the usual Tarskian models), but to a mapping P from D^* to the set TV of truth values.[20]

MVL's can supply a logical (non-probabilistic) basis for dealing with uncertainty, if we read truth values as measures of the strength of our belief in the truth of a proposition. To this respect, a particularly interesting case is when $TV = [0, 1]$; then, the semantic interpretation function (given a multi-valued model M and an assignment σ) will be as follows

$$
\begin{aligned}
[P(t_1, \ldots, t_k)]_{M,\sigma} &= P([t_1]_{M,\sigma}, \ldots, [t_k]_{M,\sigma})\\
[\sim\alpha]_{M,\sigma} &= 1 - [\alpha]_{M,\sigma}\\
[\alpha \vee \beta]_{M,\sigma} &= \max([\alpha]_{M,\sigma}, [\beta]_{M,\sigma})\\
[\alpha \wedge \beta]_{M,\sigma} &= \min([\alpha]_{M,\sigma}, [\beta]_{M,\sigma})\\
[\alpha \rightarrow \beta]_{M,\sigma} &= \min(1, 1 - [\alpha]_{M,\sigma} + [\beta]_{M,\sigma})\\
[\forall x \cdot \alpha]_{M,\sigma} &= \inf_{d \in D}([\alpha]_{M,\sigma[x/d]})\\
[\exists x \cdot \alpha]_{M,\sigma} &= \sup_{d \in D}([\alpha]_{M,\sigma[x/d]})
\end{aligned}
$$

where the interpretation of the terms and the meaning of $\sigma[x/d]$ are the usual ones. Notice that MVL's are characterized by the behaviour of the truth values with respect to the logical connectives, behaviour which is quite different from that of the certainty values in the numerical techniques seen above. We will soon see how Zadeh actually based his "fuzzy logic" on an underlying logic which is just the MVL described. As for three-valued logics, they have been used to model incomplete belief (e.g. Levesque, 1984), or to give semantics to logic programs (Barbuti-Salibra, 1986).

The modal logic LL

The second non-standard logic we consider for the uncertain knowledge representation task is a modal extension of PC due to Halpern and Rabin (1983). It is perhaps worthwhile to recall that a modal logic is one in which modal distinctions (like "necessary", "possible", etc.) can be expressed by modal operators which must take into account not the extension (truth value) of their arguments, as is the case for the usual logical operators, but their intension (the concepts they express) instead. To illustrate the difference, we notice that while the truth of

$\sim$practise(Alessandro, hara-kiri)

[19] Bochvar and Kleene later defined their own 3-valued logics, the main difference lying in the interpretation given to the third value—"undeterminate" to Lukasiewicz, "unknown" to Kleene and "meaningless" to Bochvar—and therefore in the way it behaves with respect to the logical connectives.

[20] Notice that the classical case is caught by taking TV = {true, false}.

is fully determined by the truth of "practise(Alessandro, hara-kiri)", this is not the case for the truth of

$$\mathbf{M}\ \text{practice(Alessandro, hara-kiri)} \tag{8}$$

where **M** is a modal operator to be read as "it is possible that".

Because of this intentional character of modal logics, a satisfactory semantics was not devised until 1963, when S. Kripke proposed his "possible world semantics", shifting the focus of attention from the truth of a proposition to the interpretations (possible worlds) in which that proposition holds. On these possible worlds (different states of affairs) he defines an "accessibility" relation that intuitively tells us which worlds are "possible alternatives" given a particular one. We can easily give semantics to formulas like (8) in the Kripke framework: (8) will be true (in a certain state of affairs w) if there is a possible alternative state w′ in which "practise(Alessandro, hara-kiri)" is true. It can be easily seen that different properties of a modal operator can be captured by imposing different constraints on the accessibility relation.[21]

Halpern and Rabin, resurrecting an idea from Rescher (1968), build their logic by defining a modal operator **L** with the intended meaning of "it is probable that", allowing qualitative reasoning about likelihood (vs. the quantitative one modelled by the MVL's), like

$$\frac{\mathrm{LB} \quad \mathrm{B} \rightarrow \mathrm{LA}}{\mathrm{LLA}}$$

(if B is probable, and A is probable when B is the case, then A is somewhat probable). The semantics of LL are given in a possible-world framework in which the intended meaning of the accessibilty relation (which must be reflexive) is that world w′ is "accessible" from a world w if the state of affairs it expresses is a probable one given that expressed by w; so **L**A is true in w if there is a world w′, accessible from w, in which A holds. In a later paper, Halpern and McAllester (1984) devise a complex machinery to link LL to probabilities, and enrich LL with a knowledge operator **K**. More recently, Halpern and Rabin (1987) further investigated the LL logic and extended it by allowing a modal operator for "conceivability"; examples of how LL can be used are also given in this work.

5.2 Fuzzysets and fuzzy logic

In 1965 Lotfi A. Zadeh noticed that a host of the concepts commonly used in human reasoning (like "old", "warm", "high" and so on) are intrinsically vague and subjective, and so is the boundary between the objects which fit them and those which don't. He then devised his "fuzzy sets". Later Zadeh also realized that most of the quantifiers used in natural language (e.g. "most of"), as well as the meta-linguistic statements about the truth of statements (e.g. "that's pretty true"), are vague as well. This led to his "fuzzy logic".

"Fuzzy sets" theory (Zadeh, 1965) starts from a vague interpretation of predicates, according to which a fuzzy predicate P on a domain D is a function $P{:}D \rightarrow [0, 1]$ measuring to what extent an element of D satisfies P. P defines a fuzzy set

$$Ap = \{\langle e, x\rangle \mid e \in D,\ x \in [0, 1],\ P(e) = x\}$$

where x is called the "degree of membership" of e. Conversely, given a fuzzy set A, his characteristic function (predicate) will be

$$f_A = \lambda ex \cdot x$$

Having redefined sets, what Zadeh needed was a coherent definition of the operations on them; he followed the MVL suggestion and proposed the definitions in Fig. 6, with A, B fuzzy subsets of D. Notice that fuzzy sets do not obey ordinary set-theory rules (e.g. $A \cap \sim A = \phi$ does not hold: if $\langle e, 0.6\rangle \in A$, then $\langle e, 0.4\rangle \in A \cap \sim A$) resembling the MVL's.

[21] See, for instance, (Huges-Cresswell, 1968) on modal logics; (Turner, 1984) gives a clear synthetic AI oriented presentation.

$$A \cup B = \{\langle e, x\rangle \mid e \in D,\ x = \max(f_A(e), f_B(e))\}$$
$$A \cap B = \{\langle e, x\rangle \mid e \in D,\ x = \min(f_A(e), f_B(e))\}$$
$$\sim A = \{\langle e, x\rangle \mid e \in D,\ x = 1 - f_A(e)\}$$
$$A \subseteq B \text{ iff } \forall e \in D,\ f_A(e) \leqslant f_B(e)$$

Figure 6

Zadeh took the next step ten years later. The MVL's still in his mind, he defined a "fuzzy logic" (Zadeh, 1975) to be a logic in which the truth values are fuzzy subsets of the [0, 1] interval, called "linguistic truth values" (TV's). It is clearly

$$TV \subseteq ([0, 1] \to [0, 1])$$

and so the semantics, for a language L, will be a function $[\![\,]\!]: L^* \to ([0, 1] \to [0, 1])$.

Just to give the feeling of this semantics, we note the formula for the disjunction.

$$[\![A \,\&\, B]\!] = \lambda v \cdot \min([\![A]\!]v, [\![B]\!]v)$$

Intuitively, we see how this semantics maps each formula α to a function that tells us, given a value $v \in [0, 1]$, to what extent v can be regarded as the truth value of α. The fuzzy school then proceeded by fuzzyfying quantifiers: most of the common-sense statements are, for Zadeh, 1983 "dispositions"

Q A's are B's

where Q is a fuzzy quantifier (most, few, some,...) and A and B are fuzzy sets (or relations). By fuzzyfying inferences too:

$$\begin{array}{c} x \text{ is } A \\ \underline{\langle x, y\rangle \text{ are } B} \\ y \text{ is } (A \circ B) \end{array} \qquad \begin{array}{c} Q_1 A\text{'s are } B\text{'s} \\ \underline{Q_2 C\text{'s are } D\text{'s}} \\ Q_3 E\text{'s are } F\text{'s} \end{array}$$

(where all the relations and the quantifiers are fuzzy, and $\circ$ is a fuzzy composition operator) two (examples of) fuzzy inference schemes are obtained (the second is a general syllogism recently proposed by Zadeh, 1985).

A wide debate is swarming around Zadeh's ideas. Shefe (1980) is a sharp critic: he insists that linguistic vagueness is an uncertainty about the predicate's applicability, and it must therefore be dealt with metalinguistically rather than at the object level; inferences are likewise exact, all vagueness being reducible to the uncertainty in the adequacy of the descriptions.

Another common criticism regards the subjectivity of fuzzy definitions: the definition of, say, "old" depends on the subject and on the nature of the object involved, but also on the micro-social and macro-social context ("old" in the sports circle, "old" in a tribe) and on the purpose ("old" for social security or "old" for medical considerations;[22] and the definition of the TV's for fuzzy logics is itself subjective as well.

Zadeh's ideas have found a great number of applications in AI; space precludes even a partial presentation, but the interested reader is referred to the bibliographical notes.

5.3 The non-representation of uncertainty

All the techniques seen up to now somehow represent the state of certainty and provide mechanisms to switch from one state to another. Some AI systems have been developed which are able to cope with uncertain knowledge without representing uncertainty. These systems usually take advantage of the redundancies present in the problem: they do this by using multiple knowledge sources (KS's),

[22] cf. Cohen (1985) p. 146.

or a smart control strategy, or by guessing plausible hypotheses,[23] or even by using all of these techniques together.

Ferrante (1985) gives an example of the first technique, combining the information from different KS's while considering their "characteristic errors" (depending on the KS's structure, and the environment and data the KS works with) to decide which KS is more reliable on what. MOLGEN (Stefik, 1981) uses both the second and the third technique: it faces (or rather circumvents) uncertainty by the "least commitment approach" (a decision is delayed until there is enough information to take it safely) while guessing a plausible supposition when the information is too little to go by.

Finally, these techniques are applied all together in the HEARSAY-II speech understanding system (Erman et al., 1980). Various KS's, consisting of programs for acoustic, lexical, syntactical interpretations and the like, are used in the system, interacting via a "blackboard" model. The implementation of this model allocates resources to KS's according to the certainty and the utility of the information they provide ("opportunistic scheduler"). Through this cooperation among KS's, ambiguities that originate for instance at the acoustic level like

tell Bob rings
till Bob rings
tell Bob brings
till Bob brings

can be effectively solved. In the end, some KS's are designed to guess the next word based on the part of the sentence already heard. HEARSAY-II achieved an excellent performance in a domain deeply characterized by uncertainty while not making use of any specific uncertainty management technique. (Erman et al., 1980) insists on this approach being extensible to other domains.

6 Which is the best uncertainty calculus?

Having shown a number of uncertainty treatment techniques, we can now climb a few steps and get a larger perspective: the goal will be to identify the appropriate AI approach within the uncertainty representation jungle. In this perspective, a short discussion will be attempted, followed by a partial bibliography of the vast literature in the field.

Having reached a higher vantage point, we would expect to see two groups facing each other, the uncertainty types and interpretations on the one side, and the uncertainty management techniques on the other. We have become acquainted with many of these in the above sections. Now one might feel tempted to get rid of the apparent chaos by organizing a tournament between each of the two groups, with the intention of proclaiming the best uncertainty interpretation and uncertainty management technique. We have already seen how tempting this possibility is when comparing Dempster-Shafer's theory with the Bayesian one, or when balancing the numerical and symbolical representations of uncertainty. Not wishing to enter the epistemological arena, however, this attitude is hardly in the spirit of AI.

Once we accept the variety of uncertainty types a more promising perspective from the AI view point seems to be to view the two groups as two rows of dancers, each waiting for a partner to step forward from the opposite row. Our aim would then be to match these dancers and so find the designated partner technique for the uncertainty type affecting our problem.[24] The position according to which uncertainty theories are alternative rather than rival ones, and we are allowed to choose that which best fits our needs, has been advocated before in Szolovits and Pauker (1978); Fox (1986; 1987); Garbolino (1986). As an illustration, (Fox, 1987) proposes that a sensible choice

[23] The idea of making assumptions which will possibly have to be withdrawn later is characteristic of non-monotonic reasoning. As noted in section 1.3, also non-monotonic reasoning can be seen as a technique to cope with uncertainty without explicitly dealing with it.

[24] Upon reflection, this is nothing but the extension of the usual knowledge representation choice problem to the representation of uncertainty as a kind of knowledge.

between numerical and knowledge-based methods should select the former when we have a great deal of experience relevant to the problem at hand, while our knowledge about it is poor; while selecting the latter when the reverse is true.

In order to contemplate "matching" of techniques we must first formalize each uncertainty model, in a common framework, stressing its adequacies and its drawbacks. Questions like "Which facets of uncertainty are best expressed, and which are not?", "What preconditions need to be met?", "What structure, if any, should the domain have?", "What precision does the model require, and what does it offer?", "How does the model make uncertain pieces of knowledge interact with certain ones?" and the like should be answered in such a framework. We hope we have made some contribution to answering these questions (when not already directly answered) in the above sections. Some framework of this kind has been proposed (see next section) but usually the goal has been to tidy up the uncertainty treatment techniques zoo rather than to highlight the adequacies of each of its beasts.

The next step is to choose the technique that best fits our uncertainty problem. Roughly, this will be done in three phases:

1 select those techniques which are applicable to the problem (the problem fits their preconditions);
2 verify the epistemological and computational adequacy of the selected techniques for the uncertainty at hand;
3 weigh the remaining techniques and choose one: the general context (e.g. the available computational tools, the representation used for the other knowledge types of the problem, etc.) should be taken into account.

So expressed, the choice of an adequate technique clearly reveals itself to be a meta-problem, which can be dealt with by the use of AI techniques. This is just the path suggested by Fox (1986), who proposes an AI system in which the object knowledge describes uncertainty management techniques, and weak methods are used to deal with the uncertainty in the meta-decision of which calculus to use.

Finally, an AI system which is to treat uncertainty should be designed so that any uncertainty management techniques can be easily installed in it. Unfortunately, this is not the right place to inspect this important issue in a deeper way; more insights on this point may be found in (Saffiotti, 1988), where an architecture meeting this requirement is sketched out, and a sample Prolog implementation of such a system is proposed.

Bibliographical notes

Historical discussions

A historical perspective of how uncertainty entered the epistemological debate is given in (Hacking, 1975) and (Buchanan-Shortliffe, 1975). (deFinetti, 1937) gives a taste of the subjective school. (Polya, 1954) is a rich and intriguing book (revised edition published in 1969) where the treatment of uncertainty is embedded in a philosophical and psychological context.

Bayesian techniques

The applications of Bayes' theorem (of which an excellent account can be found in Lindley, 1971) are well illustrated in AI by (Duda et al., 1976) and (Pednault et al., 1981) for the subjective Bayesian methods, and by (Buchanan-Shortliffe, 1984) for the MYCIN adventure and its developments. Other applications can be found in (White, 1985), whose approach, lying between induction and statistics, relaxes the conditional independence constraint; in (Pearl, 1986), who is oriented towards hierarchically structured domains and uses message transmission to implement the certainty updating; and in the works discussed in section 2.4 (Friedman, 1985; Ishizuka et al., 1981; Cheeseman, 1984). The last one discusses and applies the powerful Principle of Maximum Entropy, dealt with by (Shimony, 1985) from a more philosophical and critical point of view. The reader can catch a vivid glimpse of the debate around the Bayesian methods through (Lindley, 1971), (Lindley, 1987), (Shafer, 1987), (Garbolino, 1986).

Dempster-Shafer theory

The theory of Dempster-Shafer, fully expounded in (Shafer, 1976), is briefly and clearly presented in (Barnett, 1981), (Garvey et al., 1981) and (Gordon-Shortliffe, 1984). The first proposes a linear time implementation of Dempster's rule. The behaviour of Dempster's rule is analyzed in (Ginsberg, 1984) and in (Dubois-Prade, 1985). Other AI applications of the theory, seen in section 3.3, are (Lu-Stephanou, 1981) and (Strat, 1984). (Shafer, 1987) gives further insights on the theory, and compares it with the Bayesian one.

Non-numerical approaches

At the moment work on the endorsement model seems to be limited to that of Cohen and his co-workers. In addition to his original proposal (Cohen, 1985), (Cohen, 1984) and (Cohen-Sullivan, 1985) report successive developments of the model.

The debate about the use of numerical methods and possible alternative techniques can also be followed through the early (McCarthy-Hayes, 1969) and (Davis, 1979), and in (Fox, 1986) and (Fox, 1987). (Shafer, 1987) is an advocate of numerical methods. (Tversky-Kahneman, 1974) is a rich source of insights on human reasoning under uncertainty.

Fuzzy logic

Zadeh himself is a rich source: (Zadeh, 1965 and 1975) report on Fuzzy-sets and fuzzy-logic respectively, (Zadeh, 1981) presents a natural language oriented fuzzy system (PRUF), (Zadeh, 1983) clearly illustrates the use of fuzzy logic, and (Zadeh, 1985) proposes general syllogisms and fuzzy quantifiers as a means to represent common-sense knowledge. Other fuzzy-sets or fuzzy-logic based languages have been developed (leFaivre, 1976) (Umano et al., 1978–79), and many AI systems based on these languages (more or less) have been built (e.g. Bezdek, 1976, Ishizuka et al., 1981). Not wishing to list the numerous applications that have been proposed, we simply refer to the (Gaines-Kohout, 1977) bibliography and, since 1978, to the "Fuzzy Sets and Systems" journal. (Mamdani-Gaines, 1981) is a useful guide to the applications Zadeh's ideas can find (and have found). (Schefe, 1980) gives a linguistically oriented criticism of Zadeh's ideas.

Other approaches

Among the other approaches, we cite (Yeh, 1985) for the merging or splitting of uncertainties, (Kahn-Jain, 1985) for the uncertainty treatment in a distributed environment in which various agents are to co-operate. (Shastri-Feldman, 1985) for a maximum entropy approach to semantic nets, (Rollinger, 1983) for a model in which a "vague" matching is defined for the rule premises; and (Nilsson, 1986), in which certainty is allotted to "possible worlds" (states of affairs) rather than to propositions, while the updating follows a Bayesian scheme. Finally, the use of multiple knowledge-sources to reduce uncertainty (without explicitly having to deal with it) in the field of computer vision is well exemplified (in addition to the works already cited in section 5.3) by (Glicksman, 1983).

Other reviews

Some attempt to survey and structure the wide field of uncertainty treatment in AI has been done. (Thompson, 1985) has already been mentioned and mention should be made of (Dubois-Prade, 1986) and (Prade, 1983); a rich bibliography is included in the last two papers and in (Buchanan-

Shortliffe, 1984). (Tong et al., 1983) tests various uncertainty propagation techniques, while (Domotor, 1985) arranges them in a highly formalized framework. (Kanal-Lemmer, 1986) contains contributions by many active researchers in the field, and gives an effective snapshot of the state of the art. The issue of "Statistical Science" on the calculus of uncertainty in AI and expert systems (Vol. 2, No. 1, 1987, pp. 3–44) provides a vivid picture of the debate among the advocates of different numerical methods (viz. the Dempster-Shafer vs. Bayes controversy).

Acknowledgments

Paolo Garbolino, Roberto Barbuti, Maria Simi and Fabrizio Sebastiani provided useful discussions and worthy comments on early versions of this paper. Stefano Cerri actively encouraged the research reported here. Catherine Minciotti has been of unvaluable help in reviewing the drafts. I am indebted to the Editor, John Fox, for a number of improvements in the final form of the paper.

This research was supported by my mother Margherita: my final and greatest thanks are for her.

References

Adams, J.B. (1976). "Probabilistic Reasoning and Certainty Factors." *Mathematical Biosciences* **32**. Also in (Buchanan-Shortliffe, 1984), ch. 12.

Barbuti, R. and Salibra, A. (1986). "Logic Negation in Logic Programming and its Three-Valued Semantics." Internal Report, Computer Science Department, University of Pisa.

Barnett, J.A. (1981). "Computational Methods for a Mathematical Theory of Evidence." Proc. of the 7th IJCAI, Vancouver, British Columbia, Canada.

Bezdek, J.C. (1976). "Feature Selection for Binary Data—Medical Diagnosis with Fuzzy-Sets." *AFIPS Proc.* **45**, 1057–1068.

Buchanan, B.G. and Shortliffe, E.H. (1975). "A Model of Inexact Reasoning in Medicine." *Mathematical Biosciences* **23**, 351–379. Also in (Buchanan-Shortliffe, 1981), ch. 11.

Buchanan, B.G. and Shortliffe, E.H. (Ed.) (1984). Rule Based Export Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project, Addison-Wesley.

Carnap, R. (1950). *Logical Foundations of Probability*. Univ. of Chicago Press, Chicago.

McCarthy, J. and Hayes, P.J. (1969). "Some Philosophical Problems from the Standpoint of Artifical Intelligence," in Meltzer and Michie (Eds), *Machine Intelligence* **4**, Edinburgh, 163–502.

Cheeseman, P. (1984). "Learning of Expert Systems from Data." IEEE Proc. of Workshop on Principles of Knowledge-Based Systems, Denver, Colorado, 115–112.

Cohen, P.R. (1984). "Progress Report on the Theory of Endorsements: A Heuristic Approach to Reasoning About Uncertainty." IEEE Proc. of Workshop on Principles of Knowledge-Based Systems, Denver, Colorado.

Cohen, P.R. (1985). *Heuristic Reasoning about Uncertainty: An Artificial Intelligence Approach*. Pitnam, Boston.

Cohen, P.R. and Sullivan, M. (1985). "An Endorsement-Based Plan Recognition Program." Proc. of the 9th IJCAI, Los Angeles, California.

Davis, R. (1979). "Panel on Dealing With Uncertainty," in Proc. of 6th IJCAI, Tokyo, Japan.

de Finetti, B. (1937). "La prévision- seis logiques, seis sources subjectives," *Annales de l'Institut H. Poincare* **7**, 1–68.

Dempster, A.P. (1968). "A Generalization of Bayesian Inference." *Journal of Royal Statistical Society Ser. B* **30**, 205–247.

Domotor, Z. (1985). "Probability Kinematics, Conditionals, and Entropy Principles." *Synthese* **63**, 75–114.

Dubois, D. and Prade, H. (1985). "Combination and Propagation of Certainty With Belief Functions—a reexamination," in Proc. of 9th IJCAI, Los Angeles, California.

Dubois, D. and Prade, H. (1986). "A Tentative Comparison of Numerical Approximate Reasoning Methodologies", Workshop on IDSS for Plant Operation, CCE, Ispra, 1986.

Duda, R.O., Hart, P.E. and Nilsson, N.J. (1976). "Subjective Bayesian Methods for Rule-Based Inference Systems." *AFIPS Proc.* **45**, New York, 1075–1082.

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**(2), 213–253.

Le Faivre, R.A. (1976). "Procedural Representation in a Fuzzy Problem-Solving System." *AFIPS Proc.* **45**, 1069–1074.

Ferrante, R.D. (1985). "The Characteristic Error Approach to Conflict Resolution," in Proc. of 9th IJCAI, Los Angeles, California.
Fox, J. (1986). "Three Arguments for Extending the Framework of Probability," in Kanal-Lemmer, 447–458.
Fox, J. (1987). "Making decisions under the influence of Knowledge," in P. Morris (ed.), *Modelling Cognition*. John Wiley & Sons Ltd, 199–212.
Friedman, L. (1981). "Extended Plausible Inference." Proc. of the 7th IJCAI, Vancouver, British Columbia, Canada.
Gaines, B.R. and Kohout, L.J. (1977). "The Fuzzy Decade: a Bibliography of Fuzzy Systems and Closely Related Topics." *International Journal of Man-Machine Studies* **9**, 1–68.
Garbolino, P. (1986). "A Comparison of some Rules for Probabilistic Reasoning." Workshop on I.D.S.S. for Plant Operation, J.R.C., Ispra.
Garvey, T.D., Lowrance, J.D. and Fischler, M.A. (1981). "An Inference Technique for Integrating Knowledge from Disparate Sources." Proc. of the 7th IJCAI, Vancouver, British Columbia, Canada.
Ginsberg, M.L. (1984). "Non-Monotonic Reasoning Using Dempster's Rule." Proc. of 1st National Conf. on Art. Int., AAAI, Austin, Texas.
Glicksman, J. (1983). "Using Multiple Information Sources in a Computational Vision System." Proc. of 8th IJCAI, Karlsruhe, West Germany, 1078–1080.
Gordon, J. and Shortliffe, E.H. (1984). "The Dempster-Shafer Theory of Evidence," in Buchanan-Shortliffe, ch. 13.
Hacking, I. (1975). *The Emergence of Probability*. Cambridge University Press.
Halpern, J.Y. and McAllester, D.A. (1984). "Likelihood, Probability and Knowledge," in Proc. of 1st National Conf. on Art. Int., AAAI, Austin, Texas.
Halpern, J.Y. and Rabin, M.O. (1983). "A Logic to Reason About Likelihood." Proc. of the 15th Annual Symposium on the Theory of Computing, 310–319.
Halpern, J.Y. and Rabin, M.O. (1987). "A Logic to Reason about Likelihood." *Artificial Intelligence* **32**, 379–405.
Huges, G.E. and Cressell, M.J. (1968). *An Introduction to Modal Logic*, Methuen, London.
Ishizuka, M., Fu, K.S. and Yao, J.T.P. (1981). "Inexact Inference for Rule-Based damage Assessment of Existing Structures." Proc. of the 7th IJCAI, Vancouver, British Columbia, Canada.
Kanal, L.N. and Lemmer, J.F. (Eds) (1986). *Uncertainty in Artificial Intelligence*. North Holland, 1986.
Khan, N.A. and Jain, R. (1985). "Uncertainty Management in a Distributed Knowledge Based System," in Proc. of 9th IJCAI, Los Angeles, California.
Levesque, H.J. (1984). "A Logic of Implicit and Explicit Belief." National Conf. on Art. Int., AAAI, Austin, Texas, 198–202.
Lindley, D.V. (1971). *Making Decisions*. Wiley Interscience, London.
Lindley, D.V. (1987). "The Probability Approach to the Treatment of Uncertainty in Artificial Intelligence and Expert Systems." *Statistical Science* **2**, 17–24.
Lukasiewicz, H. (1920). "O Logice Trojwartosciowcj." *Ruch Filozoficzny* **5**, 169–170; also in: S. McCall, *Polish Logic 1920–1939*, Oxford University Press, Oxford, 1967.
Lukasiewicz, J. (1930). "Many-Valued Systems for Propositional Logic," in S. McCall: *Polish Logic 1920–1939*. Oxford University Press, Oxford, 1967.
Lu, S.Y. and Stephanou, H.E. (1984). "A Set-Theoretic Framework for the Processing of Uncertain Knowledge." Proc. of National Conf. on Art. Int., AAAI, Austin, Texas, 216–221.
Nilsson, N.J. (1986). "Probabilistic Logic." *Artificial Intelligence* **28**, 71–87.
Pearl, J. (1986). "On Evidential Reasoning in a Hierarchy of Hypotheses." *Artificial Intelligence* **28**, 9–15.
Pednault, E.P.D., Zucker, S.W. and Muresan, L.V. (1981). "On the Independence Assumption Underlying Subjective Bayesian Updating." *Artificial Intelligence* **16**, 213–222.
Polya, G. (1954). *Patterns of Plausible Inference*. Princeton University Press, Princeton, New Jersey.
Prade, H. (1983). "A Synthetic View of Approximate Reasoning Techniques." Proc. of 8th IJCAI, Karlsruhe, West Germany, 130–136.
Rescher, N. (1968). *Topics in Philosophical Logic*. Reidel, Dordrecht.
Rollinger, C.R. (1983). "How to Represent Evidence—Aspects of Uncertain Reasoning." Proc. of 8th IJCAI, Karlsruhe, West Germany.
Saffiotti, A. (1987). "Un sistema di inferenze non-monotone su basi di conoscenza con incertezze", doctoral dissertation, Computer Science Department, University of Pisa (in preparation).
Schefe, P. (1980). "On Foundations of Reasoning With Uncertain Facts and Vague Concepts." *Int. Journal of Man-Machine Studies* **12**, 35–62.
Shapiro, E. Y. (1983). "Logic Programs With Uncertainties: A Tool for Implementing Rule-Based Systems." Proc. of 8th IJCAI, Karlsruhe, West Germany.
Shafer, G. (1976). *A Mathematical Theory of Evidence*. Princeton University Press, Princeton, New Jersey.

Shafer, G. (1987). "Probability Judgment in Artificial Intelligence and Expert Systems", *Statistical Science* **2**, 3–16.

Shastri, L. and Feldman, J.A. (1985). "Evidential Reasoning in Semantic Networks: A Formal Theory", in Proc. of 9th IJCAI, Los Angeles, California.

Shimony, A. (1985). "The Status of the Principle of Maximum Entropy." *Synthese* **63**, 35–53.

Strat, T.M. (1984). "Continuous Belief Functions for Evidential Reasoning." Proc. of 1st National Conf. on Art. Int., AAAI, Austin, Texas, 308–313.

Thompson, T.R. (1985). "Parallel Formulation of Evidential Reasoning Theories." Proc. of 9th IJCAI, Los Angeles, California.

Tong, R.M., Shapiro, D.G., Jeffrey, S.D. and McCune, B.P. (1983). "A Comparison of Uncertain Calculi in an Expert System for Information Retrieval." Proc. of 8th IJCAI, Karlsruhe, West Germany, 194–197.

Turner, R. (1984). *Logics for Artificial Intelligence.* Ellis Horwood Ltd., Chichester.

Tversky, A. and Kahneman, D. (1974). "Judgment under Uncertainty: Heuristics and Biases." *Science* **185**, 1124–1131.

Umano, M., Mizumoto, M. and Tanaka, K. (1978). "FSTDS System: A Fuzzy-Set Manipulation System." *Information Science* **14**, 115–149.

Umano, M., Mizumoto, M. and Tanaka, K. (1979). "A System for Fuzzy Reasoning." Proc. of 6th IJCAI, Tokyo, Japan, 917–919.

van Melle, W. (1980). "A Domain-independent System that Aids in Constructing Knowledge-based Consultation Programs." Ph.D. dissertation, report STAN-CS-80-820 c HPP-80-22, Computer Science Department, Stanford University.

Weiss, S.M. and Kulikowski, C.A. (1979). "EXPERT: A System for Developing Consultation Models." Proc. of 6th IJCAI, Tokyo, Japan, 942–947.

White, A.P. (1985). "PREDICTOR: An Alternative Approach to Uncertain Inference in Expert Systems", in Proc. of 9th IJCAI, Los Angeles, California.

Yeh, A. (1985). "Flexible Data Fusion (& Fission)", in Proc. of 9th IJCAI, Los Angeles, California, 420–422.

Zadeh, L.A. (1965). 'Fuzzy sets." *Information & Control* **8**, 338–353.

Zadeh, L.A. (1975). "Fuzzy Logic and Approximate Reasoning." *Synthese* **30**, 407–428.

Zadeh, L.A. (1981). "PRUF: A Meaning Representation Language for Natural Languages," in: E. Mamdani & B. Gaines (Eds), *Fuzzy Reasoning and its Applications.* Academic Press.

Zadeh, L.A. (1983). "Common-sense Knowledge Representation Based on Fuzzy Logic." *IEEE Computer*, October 1983, 61–65.

Zadeh, L.A. (1985). "Syllogistic Reasoning as a Basis for Combining Evidence in Expert Systems." Proc. of the 9th IJCAI, Los Angeles, California.