

PARALLEL FORMULATION OF EVIDENTIAL-REASONING THEORIES

Terence R. Thompson

Artificial Intelligence Project, System Planning Corporation
1500 Wilson Boulevard, Arlington, VA 22209

ABSTRACT

There is no general consensus on how best to attack evidential-reasoning (ER) problems, particularly in expert-system applications. Several approaches have evolved, but they have their roots in diverse fields, such as statistics and philosophy, and have neither a common terminology nor a common set of assumptions. The research reported here provides two useful results. First, it structures the evidential-reasoning problem in a general paradigm robust enough to be of practical use in design and construction of expert systems. Second, it uses this paradigm to formulate five important theoretical approaches in a parallel fashion in order to identify key assumptions, similarities, and differences. The five approaches discussed are classical Bayes, convex Bayes, Dempster-Shafer, Kyburg, and possibility.

I. STRUCTURING THE PROBLEM

The handling of evidence is a central element in such expert-system applications as diagnosis, integration, and control. In most tasks, evidence accumulates over time to dynamically affect uncertainties, so that the decision preferred earlier may differ from the one preferred later. However, delaying a decision is often not feasible, since this may foreclose opportunities or increase costs. Thus, it is important to understand how accumulating evidence will affect the decision process in the face of uncertainty.

The evidential-reasoning (ER) problem may be expressed in the following way: given reports about the world, and a set of current beliefs about the world, how shall I revise my beliefs as new reports are received? Reports may range from the simple to the complex, referring to various objects and events, and may contain various uncertainties. Beliefs also range from the simple to the complex, and have a notoriously obscure structure.

The fact that several different theoretical approaches to evidential reasoning have evolved makes it difficult to formulate and answer important questions of application. For example, what are the rules for structuring the reports about the world that feed raw material into the updating schemes advocated by each theoretical approach? What are the constraints on ER that are implicit (and explicit) in application of each of the approaches?

This leads us to seek a structured paradigm broad enough to encompass the models associated with each approach. Such a paradigm can be constructed in four parts as follows:

Background Elements - This portion of the paradigm contains a definition of the domain of discourse, that is, of the world-model to which we shall apply the ER process. It also contains current knowledge of that world including, possibly, knowledge of the cost of various actions in that world. Knowledge is described in terms of belief states.

Observation Reports - This portion of the paradigm describes the structure and content of reports about the external world that are the raw material for revision of the knowledge embedded in the background.

Updating Mechanism - This portion of the paradigm describes the assumptions, rules, and algorithms used to revise knowledge upon receipt of observation reports.

Decision Mechanism - This portion of the paradigm describes the assumptions, rules, and algorithms used to choose among various courses of action given revised knowledge of the world.

We shall use this paradigm as a framework for the remainder of the paper.

II. THEORETICAL APPROACHES

We will discuss five major approaches to evidential reasoning: classical Bayes, convex Bayes, Dempster-Shafer, Kyburg, and possibility. Each will be presented separately using the structured ER paradigm described above.

A. Classical Bayes

Background Elements - The background in this approach consists of three elements: (1) an algebra of statements, (2) a probability function defined over this algebra, and (3) a utility function defined over the same algebra. The algebra defines the domain of discourse, the probability function assigns degrees of belief to elements of the domain, and the utility function provides a means of reaching decisions in the domain when coupled with the decision mechanism.

The algebra used in the classical Bayes approach consists of base elements, operators, and statements obtained by application of the operators to the base elements. The base

elements are assumed to be mutually exclusive, so the application of the disjunctive operator alone expands the base elements into the set of all possible legal statements about the domain of discourse.

For example, if there are four mutually exclusive base elements labelled "1", "2", "3", and "4", then the set of legal statements has the following members:

(null)					
(1)	(2)	(3)	(4)		
(1 v 2)	(1 v 3)	(1 v 4)	(2 v 3)	(2 v 4)	(3 v 4)
(1 v 2 v 3)	(1 v 2 v 4)	(1 v 3 v 4)	(2 v 3 v 4)		
	(1 v 2 v 3 v 4)				

In general, there will be N_n legal statements when there are n base elements, where

$$N_n = \sum_{p=0}^n \{n! [(n-p)! p!]^{-1}\} = 2^n.$$

The second major element of the background is a probability function defined over the algebra of statements that obeys the following axioms:

$$0 \leq p(x) \leq 1, \text{ and} \\ p(x \vee y) = p(x) + p(y), \text{ if mutually exclusive.}$$

In addition, the sum of the probabilities assigned to the base elements is required to be one.

The probability function assigns numbers to the legal statements based upon these axioms. For example, if the probabilities assigned to the four base elements are each 0.25, then the legal statements have the following p -values:

(null)					
0.0					
(1)	(2)	(3)	(4)		
.25	.25	.25	.25		
(1 v 2)	(1 v 3)	(1 v 4)	(2 v 3)	(2 v 4)	(3 v 4)
.50	.50	.50	.50	.50	.50
(1 v 2 v 3)	(1 v 2 v 4)	(1 v 3 v 4)	(2 v 3 v 4)		
.75	.75	.75	.75		
(1 v 2 v 3 v 4)					
1.0					

The third major element of the background is a utility function defined over the algebra of statements. This is often construed as a loss function; it gives the loss, l_{ij} , incurred when the i^{th} action is taken in the face of the state of nature corresponding to the j^{th} base element in the algebra (B2, C1, J1).

Observation Reports — The observation reports are direct assignments of new p -values to elements of the algebra of statements. That is, they assign a number to certain

propositions that may be construed as a new degree of belief in the truth-value of that proposition. The assignment of this new p -value causes a re-assignment of p -values to all other statements in the algebra via the updating mechanism.

There are several ways in which this direct assignment of new p -values may be viewed:

- (1) Each observation report could consist of the assignment of a single p -value of 1.0 to some element in the algebra of statements.
- (2) Each observation report could consist of the assignment of a single p -value in the interval $[0,1]$ to some element in the algebra.
- (3) Each observation report could consist of the assignment of two p -values in the interval $[0,1]$ to some element in the algebra. These serve as lower and upper p -values for the element.

The primary effect of these different views is upon the size of the algebra of statements. The number of statements required is largest under the first view, since we must have a single statement corresponding to each and every possible observation (value read on a meter, etc.). The other views allow us to use fewer statements, since we may map several observations onto a single statement. Ordinarily, only the first view is utilized in the classical Bayes approach.

Updating Mechanism — The classical Bayesian approach rests upon Bayes' Rule for calculating posterior probabilities of states of nature from two items: (1) prior probabilities on those states, and (2) conditional probabilities for evidence given certain states of nature. In symbolic form,

$$P(S_i|E_j) = \frac{P(S_i) P(E_j|S_i)}{\sum_k P(S_k) P(E_j|S_k)},$$

where

$P(S_i|E_j)$ is the posterior probability of state S_i given evidence E_j ,

$P(S_i)$ is the a priori probability of state S_i (i.e., before evidence is taken into account),

$P(E_j|S_i)$ is the conditional probability of E_j given state S_i .

If we have a probability or degree-of-belief distribution on the evidence, $P(E_k)$, we compute the current p -value for each state of nature from the posterior probabilities and the evidential p -values according to a conditionalization formula such as

$$P_{\text{cur}}(S_i) = \sum_k P(E_k) P(S_i|E_k),$$

where we assume that the distribution on the evidence is normalized to one. Variations on this formula are possible depending upon the structure of the algebra of statements. Note that the formula used here is compatible with the first and second interpretations of observation reports.

In terms of the algebra of statements discussed thus far, there is no necessary differentiation between categories of statements in the algebra. That is, observation reports could be received for any one of the statements.

Some writers explicitly divide the statements in the algebra into two distinct classes: hypotheses and evidence (e.g., D2). Hypotheses are often called states of nature, while evidence is often termed a measurement. In any case, the basic idea is that there is a directionality or hierarchy in the web of inference: we reason from evidence to hypotheses.

There may be an advantage in taking some form of hierarchical approach. First, the inferential relationships between statements in the algebra are made more explicit than they are in the undifferentiated algebra. Second, the computational burden associated with each updating cycle may be lessened in that the effects of an observation report are limited to portions of the hierarchy explicitly connected with the statement set that is the subject of the report.

Whether or not the approach is hierarchical, if we use the first interpretation of observation reports, the updating mechanism operates just once. We have the a priori $P(S_i)$, we receive an observation report that assigns a p-value of 1.0 to one of the evidential statements, and we calculate a new set of $P_{cur}(S_i)$.

Under the second and third interpretations of observation reports, p-values may be less than 1.0 and may therefore change over time. Updating cycles could thus continue as long as new reports are received.

Decision Mechanism — Given that the updating mechanism provides us with p-values for the states of nature, and given that the background contains a measure of utility in the form of a loss function, we can formulate the expected loss of the i^{th} action as follows:

$$EL_i = \sum_j [l_{ij} P_{cur}(S_j)] ,$$

where the summation is over the j states of nature. The general Bayesian decision function is simply to choose, whenever a decision is required, the action that corresponds to the minimum expected loss.

B. Convex Bayes

Background Elements — The background in this approach, like the classical Bayes approach, consists of three elements: (1) an algebra of statements, (2) a probability function defined over this algebra, and (3) a utility or loss function defined over the same algebra. These elements serve the same functions as in the classical approach. The algebra again consists of base elements, operators, and statements obtained by application of the operators to the base elements.

The probability function in the convex Bayes approach differs in a significant way from the function in the classical approach. Here the function is a convex set of p-functions (L1). That is, the belief state is not characterized by a single function, but by a set of functions having the property of convexity: the set contains every linear combination of any two members of the set. In general, if there are n base elements, the belief state will correspond to a domain in a space of $(n-1)$ dimensions, since the n^{th} component of the belief state can be determined if $(n-1)$ components are known.

Observation Reports — The convex Bayes approach, like the classical approach, construes the observation reports as direct assignments of new p-values to elements of the algebra of statements. This new p-value again causes a re-assignment of p-values to all other statements in the algebra via the updating mechanism.

In addition to the three interpretations of observation reports previously discussed, there is now a fourth way in which assignment of new p-values may be viewed:

- (4) Some observation reports could consist of the assignment of two or more linked bounds on the convex set of p-values. These bounds are linked in the sense that they jointly specify limits on the set.

As before, the primary effect of these different interpretations is upon the size of the algebra of statements.

Updating Mechanism — The updating mechanism in the convex Bayes approach operates much like the updating mechanism of the classical Bayes approach. The key difference is that the entire convex set of functions comprising the belief state is used, rather than a single function.

As before, we use Bayes' Theorem to obtain conditional probabilities based upon the observation reports. The formulae are similar to those in the classical Bayes approach, but each probability is now indexed: P_r is the r^{th} member of a convex set of probability functions. Each new evidential input thus induces a mapping from one convex set of p-functions to another convex set.

It would seem that the computational burden of the updating mechanism will be increased by use of the convex set of p-functions in place of a single function. However, little work has been done in actual computation of updated convex belief states, so the extent of this burden is unclear at present.

Decision Mechanism — Upper and lower probabilities for some statement in the algebra can be taken from the convex set of $P_r(S_i|E_j)$ using the technique of supporting lines, planes, or hyper-planes (L1). However, no general procedure exists to handle upper and lower bounds in a utility function.

One method of attack is to suppose that the decision indicated is the one that minimizes the expected loss as was done in the classical Bayes approach. Using the convex sets of $P_r(S_i|E_j)$ and $P_r(E_j)$, we derive upper and lower bounds on each $P_{cur}(S_j)$ so that, for each action, there are now upper and lower bounds on the expected loss. Such expected-loss intervals for

different actions will, in general, overlap. No generally accepted method for choice of actions has yet been developed, although Kyburg (K1) and Levi (L1) have explored minimax techniques.

C. Dempster-Shafer

Background Elements - The background in this approach, like the first two approaches, consists of three elements: (1) an algebra of statements, (2) a mass function defined over this algebra, and (3) a utility function defined over the same algebra. The elements serve the same purposes as before, but the utility function has received little attention in the literature. It will be required, however, in practical applications.

The mass function serves as the basic vehicle for assignment and manipulation of degrees of belief. Mass is distributed across the set of subsets of the elements of the domain of discourse, that is, over the set S of $(2 \exp 2^n)$ propositions constructed from the 2^n atoms that were in turn constructed from the n base elements.

The mass function M_1 for subset A_1 of S has the following properties:

$M_1(A_1)$ is a real number on $[0,1]$;

$M_1(\text{null set}) = 0$;

$$\sum_i M_1(A_i) = 1$$

The value of $M_1(f_i)$ is taken to be the weight of belief that is ascribed just to f_i . The f_i for which $M_1(f_i)$ is nonzero are called focal elements of M_1 . Since S is itself a member of S , $M_1(S)$ describes the weight of belief unassigned to any smaller subset of S ; this is generally termed the uncertainty.

This approach provides two measures of belief state for a given proposition Q : support (SPT) and plausibility (PLS). They are calculated as follows (S1,S2):

$$\text{SPT}_1(Q) = \sum M_1(f_i), \text{ over } \{f_i \rightarrow Q\}$$

$$\text{PLS}_1(Q) = 1 - \sum M_1(f_i), \text{ over } \{f_i \rightarrow \sim Q\}$$

$$= 1 - \text{SPT}_1(\sim Q)$$

The support for Q is thus the sum of the mass attributed to all statements that imply Q , while the plausibility of Q is one minus the support for the negation of Q . The plausibility can also be expressed as the sum of the mass attributed to all subsets of S that contain some element of Q . It follows that the plausibility of Q is always greater than or equal to the support for Q .

The belief state concerning Q can be written as an interval using $\text{SPT}(Q)$ as the lower endpoint and $\text{PLS}(Q)$ as the upper. Some authors describe this as an interval-valued probability on Q . Kyburg has shown (K2) that closed convex sets of classical probability functions can represent belief states in a fashion that includes the mass-function representation as a special case.

The background also contains means of translating observation reports into mass functions. One method is that of a mass-function distribution; this distribution provides a normalized measure of the mass to be assigned to each element of the domain in the event of each possible observation. These distributions are analogous to the class-conditional probability density functions of standard probability theory.

Observation Reports - Observation reports, at least to the extent that they are expected to mesh with mass-function distributions, consist of statements like the following:

"The brightness of object X is between 1.2 and 1.6."

"Object X is surrounded by between 2 and 6 objects of similar brightness."

"In region Y , the expectation of encountering an object of class O_1 is much higher than that of any other class."

Each type of observation report is taken to generate a separate mass function. This presents no problem as long as it is completely clear that the evidential impact of a given report can be properly assigned to particular subsets of the domain of discourse. However, how the domain of discourse is to be structured in order to guarantee this proper assignment is not a trivial matter, since we must ensure the inclusion of subsets that can serve as recipients of mass from each and every observation report that will be received in performance of a given task.

Updating Mechanism - Suppose that we have received two observation reports that have individually engendered mass functions M_1 and M_2 . We combine M_1 and M_2 to form a new mass function, M_{12} , defined over subsets of the domain of discourse. In symbolic form,

$$M_{12}(f_k) = \frac{\sum_1 M_1(f_i) M_2(f_j)}{1 - \sum_2 M_1(f_i) M_2(f_j)},$$

where the first summation, $\sum_1$, is over all f_i and f_j such that $(f_i \cap f_j) = f_k$, while the second summation, $\sum_2$, is over all f_i and f_j such that $(f_i \cap f_j) = \text{null}$.

The updating procedure assumes that a current mass function, M_1 , is available and that a new mass function, M_2 , has been presented (based upon new observations). M_1 and M_2 are combined to form M_{12} , and this is used as the current function should other new mass functions be presented.

Decision Mechanism - The type of decision mechanism compatible with the Dempster-Shafer approach is not well understood. Support and plausibility functions for each statement in the domain of discourse can be calculated based upon the current mass function. These may be used as upper and lower bounds upon the probability of each statement, but there is as yet no accepted, general mechanism for decision-making based upon these bounds.

An expected-loss construction parallel to the classical and convex Bayesian approaches can be carried out if we construe $P_{\text{cur}}(Q_j)$ as an interval bounded by $\text{SPT}(Q_j)$ and $\text{PLS}(Q_j)$. This gives rise to difficulties of overlapping intervals similar to those encountered in the convex Bayes approach.

D. Kyburg

Background Elements - The background in this approach also consists of three elements: (1) an algebra of statements, (2) a p-function defined over this algebra, and (3) a utility or loss function defined over the same algebra. The elements serve the same functions as before, but several important differences merit discussion.

First, the concept of probability embraced by this approach is epistemological. This means that probability is actually a descriptor of credibility relative to some body of knowledge. In addition, the p-value used in this approach is an interval on $[0,1]$.

Second, Kyburg uses direct inference to assign p-values to hypotheses based upon knowledge of frequencies without requiring the assignment of precise a priori p-values. He also provides formal criteria for determining which evidence is relevant to a given statistical hypothesis and which is not.

Direct inference refers to the manner in which knowledge of chances (or frequencies, or objective probabilities) influences belief states about the outcomes of trials involving chance setups. In Kyburg's approach, some portion of the algebra of statements has the status of a body of knowledge containing statements about relative frequencies of occurrence of several characteristics in various classes.

Observation Reports - Observation reports in this approach can again be construed as statements in the algebra. When coupled with appropriate knowledge of relative frequencies, they assign new p-values to other elements of the algebra.

There is just one form of observation report in this approach: each consists of the identification of the class or classes to which the observed object belongs. Knowledge of relative frequencies then determines how to assign two p-values in the interval $[0,1]$ to some element in the algebra.

Updating Mechanism - The Kyburg approach uses direct inference for updating. In some special cases, this gives results that can be obtained from Bayes' Theorem (K1).

In order to show how the principle of direct inference is applied to the updating process, we consider an example. The body of knowledge is taken to consist of the following statements:

- (1) The fraction of members of class C_1 that have property P lies in the interval $[L_1, U_1]$.
- (2) The fraction of members of C_2 that have P lies in $[L_2, U_2]$.
- (3) The fraction of members of C_{12} that have P

lies in $[L_{12}, U_{12}]$, where C_{12} is the intersection of C_1 and C_2 .

The hypothesis of interest is that an item selected from class C_{12} has property P.

In this approach, criteria known as K-relevance and K-irrelevance (L1) provide a means of determining which evidence is relevant to a given statistical hypothesis and which is not. K-irrelevance refers to a mandatory lack of impact of a given piece of information on our deliberations concerning the credibility of a certain statistical hypothesis. The information concerning C_2 in the body of knowledge is K-irrelevant if and only if the following conditions obtain:

- (1) The current body of knowledge implies that $[L_1, U_1]$ is either a subinterval of $[L_{12}, U_{12}]$ or identical to it.
- (2) The current body of knowledge implies that $[L_1, U_1]$ is either a subinterval of $[L_2, U_2]$ or identical to it.

In our example, if the information concerning C_2 is K-irrelevant, then the information concerning C_1 is the total information K-relevant to the hypothesis.

K-irrelevance is thus a formal criterion that tells us whether or not knowledge of a specific relative frequency should influence our degree of belief that a member of C_{12} has property P. Suppose we know that:

- (1) The fraction of Swedes who are Protestant lies in $[a, b]$.
- (2) The fraction of visitors to Lourdes who are Swedish lies in $[c, d]$.
- (3) The fraction of Swedish visitors to Lourdes who are also Protestant lies in $[e, f]$.

Also suppose that we wish to attach the appropriate degree of belief to the hypothesis that a particular person is a Protestant, given that he is a Swedish visitor to Lourdes. Intuitively, we suspect that the values of a, b, c, d, e, and f will influence this degree of belief. K-irrelevance formalizes this process.

Kyburg's principle of direct inference has a simple form, once the criterion of K-irrelevance has been applied to the body of knowledge. If the information concerning C_1 is the total information K-relevant to the hypothesis, then the degree of belief to be assigned to the hypothesis is just the interval $[L_1, U_1]$.

The upshot of this process is that the Kyburg approach recommends, in many cases, that different intervals of degrees of belief be embraced. This has the consequence that the evolution of p-values as evidence accumulates will follow a different trajectory through the space of belief states. That this different trajectory may have important practical impact seems

reasonable, but remains to be demonstrated in a systematic fashion.

Decision Mechanism - The Kyburg approach offers interval-valued p -functions. As has been discussed above for both the convex Bayes and Dempster-Shafer approaches, there is currently no general decision mechanism available for interval-valued p -functions.

E. Possibility

Background Elements - The background in this approach consists of three elements: (1) an algebra of statements, (2) degree-of-membership functions defined over this algebra, and (3) a set of fuzzy decision functions defined over the same algebra. The algebra defines the domain of discourse, the membership functions assign degrees of membership to elements of the domain, and the decision functions provide a means of reaching decisions in the domain.

The degree-of-membership function is defined in terms of a fuzzy set. Such a set is made up of ordered pairs that assign a degree of membership in the fuzzy set to each value of a given characteristic. The fuzzy set is then denoted by $A = \{x_i | p_i\}$, where x_i is the i th value of the characteristic and p_i is the degree of membership of x_i in the set. In Zadeh's fuzzy logic, p -values obey the following axioms (G1,Z1):

$$\begin{aligned} 0 &\leq p(x) \leq 1 \\ p(\sim x) &= 1 - p(x) \\ p(x \wedge y) &= \min[p(x), p(y)] \\ p(x \vee y) &= \max[p(x), p(y)] \\ p(x \rightarrow y) &= \min\{1, [1 - p(x) + p(y)]\} \\ p(x \equiv y) &= \min\{|1 - p(x) + p(y)|, |1 + p(x) - p(y)|\}. \end{aligned}$$

The background also contains definitions of fuzzy predicates appropriate to the domain of discourse. For example, the following statements define the fuzzy predicate "LR" (low reflectance):

- (1) "A reflectance of 0.0 is low with degree 1.0."
- (2) "A reflectance of 1.0 is low with degree 0.6."
- (3) "A reflectance of 2.0 is low with degree 0.1."

The fuzzy set would then be represented by $LR = \{0|1.0, 1|0.6, 2|0.1\}$.

Observation Reports - Observation reports in the possibility approach provide the raw material for assignment of degrees of membership. That is, they are statements that establish the value of a characteristic of an object. Application of the membership function then determines the degree of membership of that object in relevant fuzzy sets. For example, the report "reflectance of object Z is 1.0" would establish the degree of membership of Z in the class of low-reflectance objects as 0.6.

Updating Mechanism - The possibility approach combines evidence in the following fashion. Suppose we desire to classify a certain object into one of n classes, $c_1 \dots c_n$. Based upon evidence E_1 , we develop membership functions p_{11} through p_{1n} to form the fuzzy set

$$A_1 = \{c_1 | p_{11}, c_2 | p_{12}, \dots, c_n | p_{1n}\}.$$

Similarly, for evidence E_2 ,

$$A_2 = \{c_1 | p_{21}, c_2 | p_{22}, \dots, c_n | p_{2n}\}.$$

We combine k sets of evidence to obtain

$$B(k) = \{c_1 | p(k)_1, c_2 | p(k)_2, \dots, c_n | p(k)_n\},$$

where the $p(k)_1, \dots, p(k)_n$ are integrated membership functions for each of the n classes. These are obtained from

$$p(k)_j = D_{xxx}(p_{1j}, p_{2j}, \dots, p_{kj}),$$

where D_{xxx} is one of several alternative fuzzy decision functions:

$$D_{int}(p_{1j}, \dots, p_{kj}) = \min(p_{1j}, \dots, p_{kj}),$$

$$D_{pro}(p_{1j}, \dots, p_{kj}) = \prod_{i=1}^k p_{ij},$$

$$D_{con}(p_{1j}, \dots, p_{kj}) = \sum_{i=1}^k a_{ij} p_{ij}, \left(\sum_{i=1}^k a_{ij} = 1 \right)$$

Use of D_{int} suggests that E_1 and E_2 interact in a more or less independent fashion, and that the presence of a smaller p -value should be preserved. Use of D_{pro} suggests that E_1 and E_2 interact like identical, independent trials, so that repetitive observations cause marked changes in relative values of membership. Use of D_{con} suggests that E_1 and E_2 interact in a reinforcing fashion, so that membership is intermediate between the two input values.

Decision Mechanism - The decision mechanism in the possibility approach is based upon the concepts of the fuzzy goal and the fuzzy constraint. The essential idea is that decisions are determined by the confluence of goals and constraints, and that all three are expressible as fuzzy sets (B1).

For example, suppose that the domain of discourse has been constructed to allow expression of goals and constraints in the same algebra of statements, $\{S_1, \dots, S_n\}$. Then the expression of goals and constraints would be embodied in the fuzzy sets

$$G = \{S_1 | m_{1g}, S_2 | m_{2g}, \dots, S_n | m_{ng}\},$$

$$C = \{S_1 | m_{1c}, S_2 | m_{2c}, \dots, S_n | m_{nc}\}.$$

The confluence of goals and constraints is expressed by the fuzzy set

$$DEC(G, C) = \{S_1 | m_{1gc}, S_2 | m_{2gc}, \dots, S_n | m_{ngc}\},$$

where the $m_{1gc}, \dots, m_{ngc}$ are integrated membership functions for the goals and constraints. These are obtained from

$$m_{jgc} = D_{xxx}(m_{jg}, m_{jc}),$$

where D_{xxx} is one of the fuzzy decision functions discussed

above. General criteria for choice among these decision functions in expert-system applications have yet to be developed.

Once the confluence set, $DEC(G,C)$, has been constructed, the question remains as to which decision is indicated. Several procedures are followed in the literature (B1, M1). The most notable are: (1) choice of the action having the greatest DEC degree of membership, (2) choice of an action that is a mixture of all actions weighted according to their DEC degrees of membership, and (3) choice of an action that is an equal mixture of the two actions having the minimum and maximum DEC degrees of membership. Like the decision functions themselves, general criteria for choice among these decision indicators in expert-system applications have yet to be developed.

III. DISCUSSION

From this parallel formulation of approaches, we can identify several important similarities and differences:

- (1) The structure, but not necessarily the content, of the algebra of statements is similar across the approaches.
- (2) All approaches, with the exception of the possibility approach, depend on formulation of a utility or loss function to arrive at decisions.
- (3) Structures given to belief states are significantly different. They may be points, intervals, convex sets, or fuzzy sets.
- (4) Components of the updating mechanism differ significantly, but the approaches fall into three major categories: those that use Bayes' Theorem exclusively (classical and convex Bayes), those that allow its use under certain conditions (Kyburg) or use a derivative form (Dempster-Shafer), and those that do not use it at all (possibility).
- (5) Components of the decision mechanism also differ. Extension of the expected-loss technique from the case of point-valued belief states to the cases of interval-valued or convex-set belief states may be possible, but precisely how this is to be done remains an open question. The fuzzy decision rules operating on the confluence of fuzzy goals and constraints appear to be unique.

The most salient similarity in the approaches is the dependence upon an algebra of statements that sharply defines the domain of discourse. This algebra is a fixed framework within which all observations are interpreted, all updating occurs, and all decisions are made. None of the approaches discussed here explicitly addresses the issue of a dynamic algebra that adapts to changing real-world conditions.

In sum, much remains to be done to develop a clear understanding of these and other approaches to evidential reasoning. Research in our laboratory is continuing along two principal lines:

(1) The parallel formulation is being extended to include other approaches, such as Shortliffe's certainty theory and Cohen's endorsement theory.

(2) Computational experiments are being carried out in order to compare the evolution of belief states when controlled by different approaches.

This should ultimately result in guidelines that can aid us in choosing the most appropriate evidential-reasoning approach for different types of expert-system tasks.

REFERENCES

- B1 - Bellman, R., and Zadeh, L. Decision-making in a Fuzzy Environment. Technical Report ERL-69-8, Electronics Research Laboratory, University of California, Berkeley, 1969.
- B2 - Blackwell, D., and Girshick, M. Theory of Games and Statistical Decisions, Wiley, New York, 1954.
- C1 - Chernoff, H., and Moses, L. Elementary Decision Theory, Wiley, New York, 1959.
- D1 - Dempster, A.P. A generalization of Bayesian inference. J. Royal Stat. Soc., 30(2), 205-247, (1968).
- D2 - Duda, R., Hart, P., and Nilsson, N. Subjective Bayesian Methods for rule-based inference systems. In Proceedings of AFIPS 1976 National Computer Conference, pp. 1075-1082, 1976.
- G1 - Gaines, B. Fundamentals of decision: probabilistic, possibilistic, and other forms of uncertainty in decision analysis. In Fuzzy Sets and Decision Analysis, H. Zimmerman, L. Zadeh, and B. Gaines, eds., North-Holland, Amsterdam, pp. 47-65, 1984.
- J1 - Jeffrey, R. The Logic of Decision, 2nd edition, University of Chicago Press, Chicago, 1983.
- K1 - Kyburg, H.E., Jr. The Logical Foundations of Statistical Inference, Reidel, Boston, 1974.
- K2 - Kyburg, H.E., Jr. Bayesian and non-Bayesian evidential updating. Technical Report 139, University of Rochester Computer Science Department, 1984.
- L1 - Levi, I. The Enterprise of Knowledge, MIT Press, Cambridge, Mass., 1980.
- M1 - Mamdani, E., Ostergaard, J., and Lembessis, E. Use of fuzzy logic for implementing rule-based control of industrial processes. In Fuzzy Sets and Decision Analysis, H. Zimmerman, L. Zadeh, and B. Gaines, eds., North-Holland, Amsterdam, pp. 429-445, 1984.
- S1 - Shafer, G. A Mathematical Theory of Evidence, Princeton University Press, Princeton, N.J., 1976.
- S2 - Strat, T. Continuous belief functions for evidential reasoning. In Proceedings of the National Conference on Artificial Intelligence, pp. 308-313, 1984.
- T1 - Thompson, T. Evidential reasoning in expert systems for image analysis. Technical Report 1068, Artificial Intelligence Project, System Planning Corporation, 1985.
- Z1 - Zadeh, L. A theory of approximate reasoning. Machine Intelligence 9, pp. 149-194, J. Hayes, D. Michie, and L. Kulich, eds., Wiley, New York, 1979.