EPISTEMIC LOGICS, PROBABILITY, AND THE CALCULUS OF EVIDENCE

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

ABSTRACT

This paper presents results of the application to epistemic logic structures of the method proposed by Carnap for the development of logical foundations of probability theory. These results, which provide firm conceptual bases for the Dempster-Shafer calculus of evidence, are derived by exclusively using basic concepts from probability and modal logic theories, without resorting to any other theoretical notions or structures.

A form of epistemic logic (equivalent in power to the modal system S5), is used to define a space of *possible worlds* or *states of affairs*. This space, called the *epistemic universe*, consists of all possible combined descriptions of the state of the real world and of the state of knowledge that certain rational agents have about it. These representations generalize those derived by Carnap, which were confined exclusively to descriptions of possible states of the real world.

Probabilities defined on certain classes of sets of this universe, representing different states of knowledge about the world, have the properties of the major functions of the Dempster-Shafer calculus of evidence: *belief functions* and *mass assignments*. The importance of these *epistemic probabilities* lies in their ability to represent the effect of uncertain evidence in the states of knowledge of rational agents. Furthermore, if an epistemic probability is extended to a probability function defined over subsets of the epistemic universe that represent true states of the real world, then any such extension must satisfy the well-known interval bounds derived from the Dempster-Shafer theory.

Application of this logic-based approach to problems of knowledge integration results in a general expression, called the *additive combination formula*, which can be applied to a wide variety of problems of integration of dependent and independent knowledge. Under assumptions of probabilistic independence this formula is equivalent to Dempster's rule of combination.

The research work described in this paper was sponsored by the U.S. Army Signal Warfare Center, under Contract DAAL02-85-C-0082, and by the National Science Foundation, under Grant DCR-85-13139.

I Introduction

The research work presented here was motivated by the need to improve the understanding of issues in the analysis and interpretation of evidence. In the context of this paper, the term *evidence* is used to describe the information, usually imprecise and uncertain, that is conveyed by observations and measurements of real-world systems. We have sought to gain such an understanding by examining the basic concepts, structures, and ideas relevant to the characterization of imprecise and uncertain knowledge.

Our approach is strongly based on Carnap's methodology (Carnap 1956, 1962) for the development of logical foundations of probability theory. In his formulation, Carnap developed an *universe* of possible worlds that encompasses all possible valid states of a real-world system. Information about that system, if *precise* and *certain*, identifies its actual state (e.g., a detailed diagnosis of a disease). If *imprecise* but certain, this information identifies a subset of possible system states (e.g., a number of possible diagnoses). If *uncertain*, then the information induces a probability distribution over system states (e.g., probability values for specific diagnoses).

It is important to note, however, that in Carnap's characterization no distinction is drawn between *degrees of precision* or *detail* when the information is uncertain. This representational shortcoming renders impossible the modeling of information that only assigns degrees of likelihood values to some subsets of possible states (i.e., instead of prescribing those values over all such subsets that are of relevance to the modeler). This type of information, providing some knowledge about the underlying probability distributions but not all the distribution values, is quite common in practical applications (e.g., in a medical diagnosis problem, tests and existing medical knowledge indicate that there is a 60% chance of liver disease but fail to provide any information about the likelihood of individual instances thereof).

Seeking to generalize Carnap's approach to allow for the treatment of this type of uncertain information we directed our attention to epistemic logics — a form of modal logics developed to deal with problems of representation and manipulation of the states of knowledge of rational

agents. Originally studied by Hintikka (Hintikka 1962), their use in artificial intelligence problems was proposed by Moore (Moore 1980). Recently epistemic logics have also been applied to the design of intelligent robots (Rosenschein and Kaelbling 1986).

In our extension of the Carnapian ideas the starting point is a generalization of Carnap's space of possible worlds, or *universe*. This generalization, obtained by considering representations of both the state of the world and the knowledge of rational agents, is called the *epistemic universe*. Described in the next section, the epistemic universe contains several interesting and important subset families. Two of these collections have as members *truth sets* and *support sets*, which are related, respectively, to different ontological and epistemological properties of possible worlds. Furthermore, these families have the properties of sigma algebras, i.e. the basic domain of definition of probability functions.

Again following Carnap's lead we define probabilities on these sigma algebras and consider their relationships. We differ from Carnap, however, in that we view evidence as generally providing information about the truth of some propositions while failing to give any indication about the truth of others. Evidence is further regarded as a potential modifier of our state of knowledge; accordingly, uncertain evidence is represented as a conventional[1] probability function defined on the algebra of epistemic sets. This probability is then shown to have the structure of the basic functions of the Dempster-Shafer calculus of evidence (Dempster 1967, Shafer 1976). Furthermore, if such an *epistemic probability* is extended to the sigma algebra of the truth sets (representing probabilities of the truth of propositions that describe the world), then the extension must satisfy the bounds of the Dempster-Shafer theory. These bounds correspond to the well-known concepts of lower and upper probability functions and, in this particular regard, our results are in agreement with the characterization made by Suppes (Suppes 1974) of the role of uncertain information in determining the probability distribution values that underlie rational choices in decision problems.

Our approach is also related in several ways to the probabilistic logic approach of Nilsson (Nilsson 1986) — the major differences being in the use of epistemic concepts and the derivation of global conditions for probability extension, in contrast to formulas derived from interval probability theory or from approximate-estimation techniques.

In addition, this work has similarities with that of Halpern and McAllester (Halpern and McAllester 1984) — the dissimilarities in this case being in the methods used to model uncertainty. It is important to note, however, that Halpern and McAllester represent likelihood formally as the probability of knowledge (in the epistemic-logic sense) of propositional truth, using an interpretation that is similar to ours in several significant respects.

Section 4 deals with the problems associated with the combination of the knowledge of several mutually trusting agents. Under assumptions that guarantee that the integrated knowledge is solely the logical consequence of the states of knowledge of the agents, several results are presented, including a general formula for knowledge combination. This *additive combination formula* may be applied to several knowledge integration problems involving either dependent or independent evidential bodies. For the latter case, the corresponding result generalizes the Dempster's rule of combination.

It is important to emphasize that the results of Sections 3 and 4, identifying the Dempster-Shafer calculus of evidence with the probability calculus in the epistemic universe, were derived by the direct application of conventional probability theory concepts without having to introduce other multivalued logic notions. The insight gained by using an epistemic model as the basic foundation of the Dempster-Shafer calculus of evidence has made possible the extension of this evidential formalism by the incorporation of new formulas for combining dependent evidence and for utilizing conditional knowledge.

In the exposition that follows, we have not included the proofs of any of the theoretical results obtained in the research being discussed, as such extensive discussion is well outside the scope of this paper. The reader interested in the actual details will find them discussed in a related work (Ruspini 1986a).

II The Epistemic Universe

A. The Carnapian Universe

Carnap's logical approach to probability starts with the construction of a space of possible worlds that encompasses all valid states of a system of interest. First, all propositions (actually instantiated first-order-logic predicates in Carnap's formulation) of relevance to the system

$$p, q, r, s, \ldots ,$$

are considered. All possible conjunctions of the type

$$p \wedge \neg q \wedge \neg r \wedge s \wedge \ldots ,$$

where every proposition appears only once either as itself or as its negation, are then considered. After discarding logical impossibilities, the resulting set of logical expressions includes all possible system states that may be represented using the propositions $p, q, \ldots$.

Each such state corresponds to the truth of an *atomic* proposition about the system in question. These atomic propositions are equivalent to the *elementary events* introduced in most treatments of basic probability theory.

[1] This function, satisfying the axioms of probability theory, is defined unambiguously on the sigma algebra of epistemic sets.

Obviously, by construction, only one such proposition can truly describe the state of the world. The space of atomic propositions, or *universe*, is therefore a collection of all possible alternative states of the system.

Possible worlds can also be regarded as functions that map each relevant proposition into its truth-value (i.e. **true** or **false**) or, alternatively, as subsets of true propositions (i.e., those mapped into the **true** truth-value). If a possible world is viewed through a "conceptual microscope" as illustrated in Figure 1, it can be seen to contain all true propositions in that world, including the negations of those that are false. Two possible worlds will always be different since at least one proposition which is true in one of them will be false in the other.

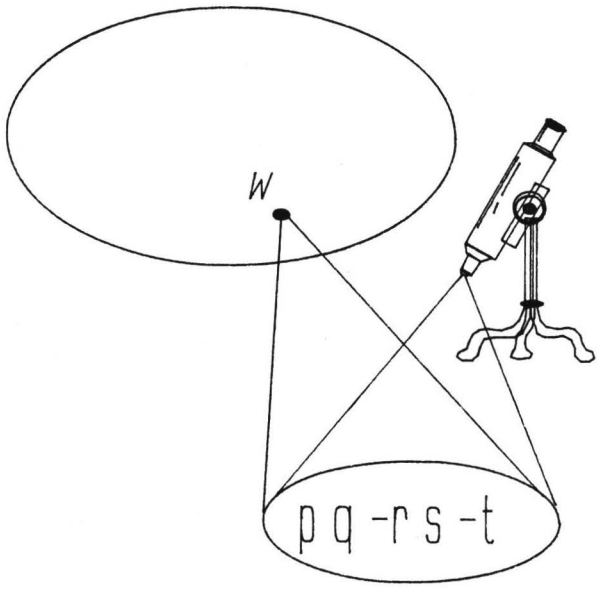

Figure 1: The Carnapian Universe Under the Microscope.

The space of possible worlds (considered as a probabilistic space) is the basic structure used by Carnap to relate the values of probability functions of subsets associated with relevant propositions on the basis of the logical relationships between those propositions.

B. Epistemic Considerations

Carnap's logical approach, while enabling a clearer understanding of the relations between logical and probabilistic concepts, suffers from a major handicap: it assumes that observations of the real world always determine unambiguously probability values for every subset in the universe. This assumption leads inevitably to problems associated with the need to define probability values when the underlying information is not rich enough to furnish them.

If, for example, we have *certain* (i.e., sure) information that a guest to a party we are hosting is fond of French wine, we would ordinarily consider, in a nonprobabilistic setting, that this information constrains our spectrum of beverage choices (assuming, of course, that we aim to please our guest and are able to do so) without identifying what particular label or vintage he is likely to prefer. If, instead of being sure, our informant is uncertain and believes there is an 80% chance that our caller will like French wine and a 20% chance that he will opt for beer, it is unreasonable (simply because uncertainty has now entered the picture) to assume that this information can be used to assign probabilities for particular choices of wine or beer when before, in a world of certainty, we regarded similar information as being only capable of identifying a subset of possibilities.

These considerations have led to the development of schemes to represent uncertain information as constraints on the values of valid probability distributions. Interval probability theories (Williams 1976), of which the Dempster-Shafer calculus is a particular case, are important examples of this technique.

The approach we have followed here, however, proceds from a different logical foundation. Starting from the notion that certain information improves our knowledge by reducing the scope of possible valid states, it considers that uncertain information is associated with a probability function defined on some subsets (actually a sigma algebra) of the universe, rather than on *every* subset of the universe. While in the case of certain information we say that *we know* that the system state is in a subset of possible states, in the case of uncertain information we similarly affirm, *with some degree of likelihood,* that that state is in in a certain region of the universe. The corresponding probability values constrain the values of other probability functions defined over richer subset collections (i.e., probability extensions).

To identify a model that constitutes the basis for defining probabilities that take values over epistemic structures, we must look at abstract formalisms that allow proper differentiation between states of the world and states of knowledge. This framework is provided by epistemic logics.

C. Epistemic Logics and Epistemic Universes

The starting point for our generalization of the Carnapian universe is again a collection of propositions about the real world, denoted by $p, q, r, s, \ldots$. We consider, in addition, more complex propositions obtained therefrom by negation, conjunction, and disjunction.[2] The resulting set of propositions is called a *frame of discernment.* Each of its members, describing a state of the world, is called an *objective proposition* or *objective sentence.*

In addition to objective sentences, we shall also deal

[2]We also consider the material implication $\rightarrow$ defined, as usual, in terms of the negation and disjunction operations.

with propositions that represent states of knowledge about the real world. When only one rational agent is concerned, the simplest of these *epistemic* propositions are denoted by $\mathbf{K}p, \mathbf{K}q, \mathbf{K}r, \ldots$, representing knowledge of their corresponding objective counterparts. We shall also consider expressions formed by combination of epistemic and objective propositions through disjunction, conjunction, implication, and negation, as in the examples $\neg\mathbf{K}r$, or $p \vee \mathbf{K}(q \vee \mathbf{K}s)$. The set of all such propositions, which encompasses the frame of discernment as a subset, is called the *sentence space*, denoted by $\mathcal{S}$.

The next step in constructing an extension of the Carnapian universe is the generation of all possible states by the assignment of truth-values to propositions in the sentence space. In addition to compliance with the axioms of ordinary propositional logic, we shall also need the following axioms, which supply the unary operator $\mathbf{K}$ with the required epistemilogical semantics:

(E1) If $\mathbf{K}p$ is true, then p is true.

(E2) If $\mathbf{K}p$ is true, then $\mathbf{KK}p$ is true (*positive introspection*).

(E3) If $\mathbf{K}(p \rightarrow q)$ is true, then $\mathbf{K}p \rightarrow \mathbf{K}q$ is also true.

(E4) If $\neg\mathbf{K}p$ is true, then $\mathbf{K}\neg\mathbf{K}p$ is true (*negative introspection*).

(E5) If p is an axiom, then $\mathbf{K}p$ is an axiom.

This system is equivalent to the modal logic system S5 (Hughes and Creswell 1968).

The space of possible worlds generated on the basis of the above schemata is called the *epistemic universe* and is denoted by $\mathcal{U}(\mathcal{S})$. When seen through our imaginary conceptual microscope, as shown in Figure 2, each possible world includes, as before, all objective propositions that are true in that world. Each possible world, however, includes also all true epistemic propositions representing knowledge of the truth (e.g., $\mathbf{K}p$) or falsehood ($\mathbf{K}\neg p$) of propositions and, in addition, propositions describing ignorance regarding the truth or falsehood of certain propositions (e.g., $\neg\mathbf{K}p \wedge \neg\mathbf{K}\neg p$).

It is important to note that, in the epistemic universe, possible worlds may share the same set of true objective propositions, even though the states of knowledge (i.e., true epistemic propositions) will be different in each case.

In the remainder of this work we will require to employ two important relations.

The first, called *logical implication* and denoted by $\Rightarrow$, holds between propositions in sentence space. This relation, well known in modal logic, is used to indicate the fact that in any possible world the truth of some proposition implies that of another. In other words, if $p \Rightarrow q$, then it is logically impossible for q to be false if p is true.

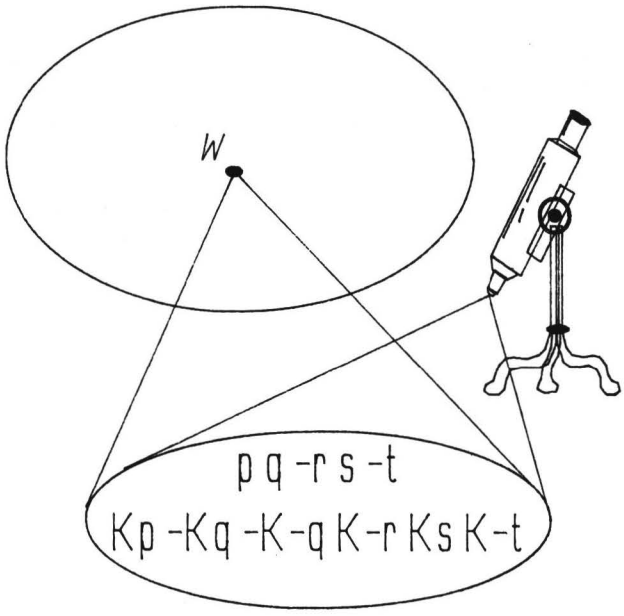

Figure 2: The Epistemic Universe Under the Microscope.

The second relation, called the *accessibility relation* and denoted by $\sim$, holds between possible worlds in the epistemic universe. Two possible worlds are related through the accessibility relation if the same epistemic propositions are true in both worlds.[3] Clearly, such world pairs cannnot be discriminated on the basis of the information (i.e. knowledge) available in each of them.

D. Special Sets in the Epistemic Universe

Several subsets of the epistemic universe are of importance in the definition of probability functions that adequately represent the effects of uncertain evidence in knowledge states.

The subset of all possible worlds where an objective proposition p is known to be true, i.e. in which the epistemic sentence $\mathbf{K}p$ is true, is called the *support set* of p and is denoted by $\mathbf{k}(p)$.

The *epistemic set* for an objective proposition p is the set of all possible worlds in which p is the *most specific proposition* that is known to be true (i.e., p is the conjunction of all objective propositions q such that $\mathbf{K}q$ is true). The epistemic set $\mathbf{e}(p)$ consists of possible worlds where $\mathbf{K}q$ is true if and only if q is logically implied by p, i.e., $p \Rightarrow q$. Pairs of possible worlds in the same epistemic set are always related by the accessibility relation $\sim$.

Epistemic sets and support sets are related by the set

[3] As characterized here, this relation is identical to the accessibility relation of modal logic only in the case of systems such as S5, where the accessibility relation is assumed to have the properties of a classical equivalence relation.

equation

$$\mathbf{k}(p) = \bigcup_{q \Rightarrow p} \mathbf{e}(q)$$

which is of essential importance to establish the relationship between epistemic constructs and the Dempster-Shafer calculus. Epistemic sets corresponding to different propositions (i.e., those that are not logically equivalent — denoted simply by $\neq$ in this work) are disjoint — the above expression, therefore, represents the disjoint partition of support sets in terms of epistemic sets. Furthermore it can be proved that

$$\mathbf{e}(p) = \mathbf{k}(p) \cap \bigcup_{\substack{q \Rightarrow p \\ q \neq p}} [\overline{\mathbf{k}(q)}] .$$

Finally, *truth sets* are important subsets of the epistemic universe that are directly related to the truth of objective, rather than epistemic, propositions. The truth set $\mathbf{t}(p)$ for an objective proposition p is the collection of all possible worlds where the prposition p is true.

Since p is true in a possible world $\mathcal{W}$ whenever $\mathbf{K}p$ is true in $\mathcal{W}$, then it follows that the support set $\mathbf{k}(p)$ is a subset of the truth set $\mathbf{t}(p)$. It is also true that $\mathbf{k}(p)$ is the largest support set contained in $\mathbf{t}(p)$.[4]

The inclusion relations between truth, support and epistemic sets are graphically illustrated in Figure 3. This figure shows the truth set $\mathbf{t}(p)$ for a proposition p; its corresponding support set $\mathbf{k}(p)$; and the epistemic sets for several propositions which imply p (including the epistemic set for p itself). As noted before, epistemic sets $\mathbf{e}(q)$ for propositions q that *do not imply* p are disjoint from the support set $\mathbf{k}(p)$ and intersect the complement $\overline{\mathbf{t}(p)}$ of the truth-set $\mathbf{t}(p)$.

III Epistemic Probabilities

A. Sigma Algebras

The collections of subsets defined in the previous section are of particular importance in a number of respects.

First, epistemic and support sets have a clear epistemological interpretation as representations of similar states of certain (i.e., sure) knowledge. Furthermore, the effect of uncertain information on states of knowledge can be represented by probability values assigned to these sets.

Truth sets, on the other hand, represent states of the world that share some ontological property. Probability values assigned to these sets represent the likelihood of certain events in the real world — namely, the truth of the proposition associated with the truth set. Because of the relations between knowledge and truth embodied in the axiom schema (E), these probability values can be expected to bear some relation, however, to probability values over support and epistemic sets. This relationship is discussed below.

Truth sets, on one hand, and epistemic and support sets, on the other, generate (by union, intersection, and complementation) *sigma algebras* of the epistemic universe, called the *truth algebra* and the *epistemic algebra*, respectively. Sigma algebras[5] are the proper domain of definition for probability functions. This fact has often been ignored in the past when, usually for the sake of simplicity, probabilities have been assumed to be defined on *every* subset of some space. Consideration of the proper domain of definition for probabilities is, however, a most important issue in probability theory (e.g., when relating joint and marginal distributions).

B. Probabilities, Supports and Masses

A probability function defined over the sigma algebra of support and epistemic sets is called an *epistemic probability*. Epistemic probabilities represent the effect of uncertain evidence on a rationl agent's state of knowledge. This effect can always be represented without ambiguity as the result of either previous experience or rational considerations. Under conditions of perfect probabilistic information (in conventional approaches this is assumed to be *always* available) the corresponding probability is defined for each atomic proposition. At the opposite end, the *vacuous* epistemic probability function assigns a probability of 1 to the epistemic set $\mathbf{e}(\mathcal{U})$ and a probability of 0 to every other subset (i.e., the evidence does not convey any information).

Two functions, both defined in the frame of discernment, can be associated in a natural manner with an epistemic probability.

The first of these, called a *mass function* and denoted by m, is defined by the expression

$$m(p) = \mathbf{P}(\mathbf{e}(p)) ,$$

i.e., as the probability of the epistemic set associated with the objective proposition p.

The second function is called the *support function* and is denoted by S. It is defined by the expression

$$S(p) = \mathbf{P}(\mathbf{k}(p)) .$$

Support functions and mass functions are related by the equation

$$S(p) = \sum_{q \Rightarrow p} m(q) ,$$

which is valid for every objective proposition p in the frame of discernment. From this basic equation, by using results from combinatorial theory (Hall 1986), it is possible

[4] Furthermore, use of the negative introspective axiom (E5) shows that that $\mathbf{k}(p)$ is the largest union of epistemic sets that is contained in $\mathbf{t}(p)$. This axiom, however, is not required to prove any of the other results discussed in this work.

[5] Sigma algebras are families of subsets that are closed under the set-theoretic operations of complementation, countable union, and countable intersection.

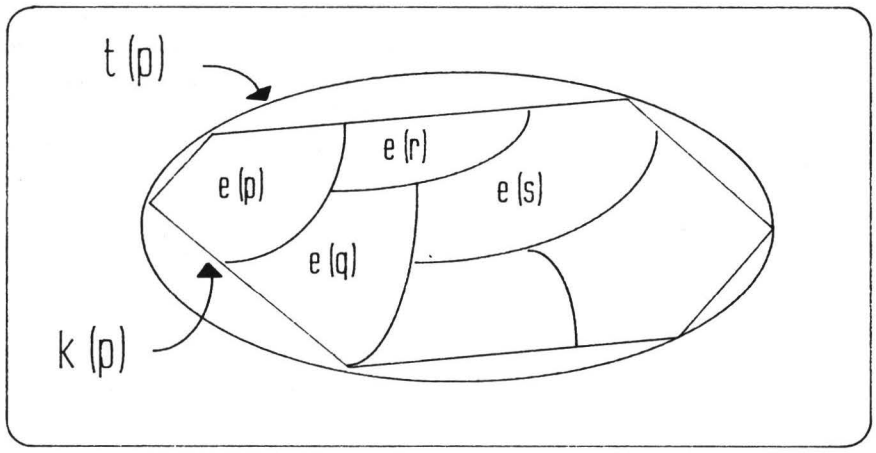

Figure 3: Relations Between Epistemic, Support, and Truth Sets.

to show that S and m are belief and mass functions, respectively, in the sense of Shafer (Shafer 1976).

In particular, it may be seen that m is expressed in terms of values of the support function S by the equation

$$m(p) = \sum_{q \Rightarrow p} (-1)^{|q-p|} S(q)\,,$$

where $|p-q|$ is the number of different (i.e., not logically equivalent) propositions r such that $q \Rightarrow r \Rightarrow p$, and where the sum is over all propositions q that imply p.

Furthermore, the following inequality, utilized by Shafer as an axiom for belief functions, can be derived as a necessary and sufficient condition characterizing support functions:

$$S(p_1 \vee \cdots \vee p_n) \geq \sum_{\substack{I \subseteq \{1,\cdots,n\} \\ I \neq \emptyset}} (-1)^{|I|+1} S\Big(\bigwedge_{i \in I} p_i\Big)\,,$$

where $|I|$ is the cardinality of the index subset I.

It is important to emphasize that the epistemic probability $\mathbf{P}$ associated with mass and support functions is a conventional probability defined on the epistemic algebra of the epistemic universe.

C. Lower and Upper Probabilities

Since both truth sets and epistemic sets are subsets of the epistemic universe, it is reasonable to ask what kind of relations exists between the probability values of members of either class. Answers to this question are obtained by considering the problems associated with the extension of an epistemic probability to a probability function defined over the truth algebra.

The problem of probability extension has received a great deal of attention in probability theory (see, for example, Neveu, 1964). The standard procedure for its solution is to define lower and upper probabilities for sets not included in the domain of definition (i.e., sigma algebra) of the probability function being extended.

The lower probability of a set X is the probability of the largest subset[6] of the sigma algebra (i.e., where the probability is actually defined) contained in X. Similarly, the upper probability of X is the probability of the smallest measurable subset that contains X.

If $\mathbf{P}_*$ and $\mathbf{P}^*$ denote the lower and upper probability functions, respectively, then well-known results of probability theory state that probability extensions $\mathbf{P}$ always exist and that the value $\mathbf{P}(X)$ satisfies the inequality constraints

$$\mathbf{P}_*(X) \leq \mathbf{P}(X) \leq \mathbf{P}^*(X)\,.$$

In addition, the bounds provided by $\mathbf{P}_*$ and $\mathbf{P}^*$ may always be attained by some extension and are therefore the best possible.

If these basic theoretical results are applied to the epistemic universe, it can be seen that the value $\mathbf{P}(\mathbf{t}(p))$ of any epistemic probability extension $\mathbf{P}$ on the truth set $\mathbf{t}(p)$ must satisfy the inequality

$$S(p) \leq \mathbf{P}(\mathbf{t}(p)) \leq Pl(p)\,,$$

where Pl is the *plausibility* function of Shafer, defined by

$$Pl(p) = 1 - S(\neg p) = \mathbf{P}(\overline{\mathbf{k}(\neg p)})\,.$$

These basic results confirms the validity of the well-known interval bounds of the Dempster-Shafer calculus.

[6] More correctly, this value is equal to the probability of any maximal measurable subsets contained in the set X under consideration.

Furthermore, lower and upper probabilities provide a general methodology to assess the impact of evidence upon understanding of the real-world state. The basic approach, according to these results, consists of representing knowledge as probabilities in an appropriate epistemic algebra, followed by estimation of the values of the lower and upper probabilities of truth sets.

IV Combination of Knowledge

This section briefly describes the results of investigations concerning the combination of the uncertain knowledge of several rational agents. For the sake of simplicity the results presented here are confined to problems involving the combination of the knowledge of two agents (Extensions to an arbitrary number of agents being straightforward).

Each of these two agents is assumed to have obtained information about the state of the world through observation devices that may possibly be dependent or correlated to some degree.

Construction of the epistemic universe that includes both the possible states of knowledge of the two agents, as well as the results of their integration requires the introduction of three unary operators: $\mathbf{K}_1$ and $\mathbf{K}_2$, representing the knowledge of each agent, and the unsubscripted operator $\mathbf{K}$, describing results of knowledge combination. It is assumed that neither agent has information about the extent or nature of the information available to the other (i.e., propositions such as $\mathbf{K}_1\mathbf{K}_2 p$ are always false), and that each agent's domain of knowledge (i.e., the sentence spaces S_1 and S_2 and their related frames of discernment) may be different.

Since the operator $\mathbf{K}$ describes the results of integrating the knowledge of two agents, it is necessary to introduce an axiom that assures that the combined knowledge is solely a function of the states of knowledge being fused:

(CK1) *The proposition* $\mathbf{K}p$ *is true if and only if there exist propositions* p_1 *and* p_2 *such that* $\mathbf{K}_1 p_1$ *and* $\mathbf{K}_2 p_2$ *are true and* $p_1 \wedge p_2 \Rightarrow p$.

The epistemic universe constructed with this augmented framework is called a *logical product universe*. In this universe it is possible, as before, to define epistemic, support, and truth sets. However, since three epistemic operators are involved, these sets must be distinguished by subscripts that identify the respective knowledge sources.

If $\mathbf{e}(p), \mathbf{e}_1(p), \mathbf{e}_2(p)$, denote the epistemic sets for the proposition p that are associated with the epistemic operators $\mathbf{K}, \mathbf{K}_1$, and $\mathbf{K}_2$, respectively, then the basic set equation that relates these sets is

$$\mathbf{e}(p) = \bigcup_{p_1 \wedge p_2 = p} [\mathbf{e}_1(p_1) \cap \mathbf{e}_2(p_2)] ,$$

where the union is over propositions p_1 and p_2 (in the respective domains of knowledge of $\mathbf{K}_1$ and $\mathbf{K}_2$) such that the conjunction $p_1 \wedge p_2$ is logically equivalent to p.

If $\mathbf{P}$ is an epistemic probability in the logical universe,[7] the above set equation may be combined with basic probability results relating marginal and joint probability distributions to derive the following general expression for knowledge combination, called the *additive combination formula*:

$$m(p) = \kappa \sum_{p_1 \wedge p_2 = p} \mathbf{P}(\mathbf{e}_1(p_1) \cap \mathbf{e}_2(p_2)) ,$$

where κ is a constant that makes $\sum m(p) = 1$.

Under assumptions of independence of the *(marginal) epistemic algebras* for $\mathbf{K}_1$ and $\mathbf{K}_2$, the above formula becomes a generalization of the Dempster's rule of combination

$$m(p) = \kappa \sum_{p_1 \wedge p_2} m_1(p_1) m_2(p_2) .$$

Simple cases of combination of dependent evidence, such as those governed by *compatibility relations*, may also be derived directly from the additive combination formula, as we have discussed elsewhere (Ruspini 1986a).

In more general cases, the corresponding expressions must combine the knowledge of the two agents (expressed by the additive combination formula) with knowledge about the dependence relations between the two evidential bodies. The latter information is typically modeled as probabilities defined on a subalgebra of the epistemic algebra.

V Conclusion

This paper has presented results that closely relate probability functions in epistemic universes to the concepts and constructs of the Dempster-Shafer calculus of evidence. The epistemic structures presented above also furnish important insight that is very useful to enhance the calculus of evidence by the development of expressions that allow for different types of dependent evidence to be combined. These expressions are the current object of our investigations, which focus particularly on the problems of combining multiple evidential bodies that share common information.

In addition, we are also concerned with problems related to the use of conditional evidence (i.e., evidence that is valid only when some proposition is true). This research expands upon and enhances our previous results in this area (Ruspini 1986b).

Our long term objectives include the treatment of problems involving combination of the knowledge of multiple agents that are aware, to different extents, of the information available to one another. The corresponding issues are

[7] More precisely, this is the case if $\mathbf{P}$ is defined in the richer *product epistemic algebra* generated by the intersections $\mathbf{e}_1(p_1) \cap \mathbf{e}_2(p_2)$.

of central importance in the design of distributed artificial intelligence systems with planning and counterplanning capabilities.

ACKNOWLEDGMENTS

The author benefitted from advice and suggestions received from colleagues in the Artificial Intelligence Center. Bob Moore, Stan Rosenschein, and David Israel suggested investigation of the Carnapian viewpoint as a prospective avenue to gain a better understanding of the relationships among approaches to approximate reasoning.

John Lowrance, Tom Garvey, Tom Strat, Len Wesley, and Stephen Lesh comprised an abundant source of insight and encouragement.

I have also profited from conversations with Andy Hanson, Haim Gaifman, Pascal Fua, Leslie Kaelbling, Martin Fischler, Ray Perrault, Kurt Konolige, Piero Bonissone, Lotfi Zadeh, Terry Cronin, Ken Laws, Lorna Shinkle, Mike Georgeff, Martha Pollack, Ivan Leclerc, Todd Davies, and Oscar Firschein. To all of them, many thanks.

REFERENCES

1 Hintikka,J., *Knowledge and Belief,* Ithaca, New York: Cornell University Press, New York, 1962.

2 Moore, R., "Reasoning about Knowledge and Action," Technical Note 191, SRI International, Menlo Park, California, 1980.

3 Carnap, R., *Meaning and Necessity,* Chicago, Illinois: University of Chicago Press, 1956.

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

5 Ruspini, E.H., "The Logical Foundations of Evidential Reasoning," Technical Note 408, SRI International, Menlo Park, California, 1986.

6 Ruspini, E.H., "Approximate Deduction in Single Evidential Bodies," In *Proc. 1986 AAAI Uncertainty Workshop,* AAAI, Menlo Park, California, 1986.

7 Suppes, P., "The Measurement of Belief," *Jrnl. Royal Statist. Soc., Series B,* **36**, (1974) 160–191.

8 Dempster, A., "Upper and Lower Probabilities Induced by a Multivalued Mapping," *Annals of Mathematical Statistics,* **38** (1967) 325–339.

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

10 Hall, M. Jr, *Combinatorial Theory,* Second Edition, New York: John Wiley and Sons, 1986.

11 Neveu, J., *Bases Mathématiques du Calcul des Probabilités,* Paris, France: Masson, 1964.

12 Rosenschein, S.J. and L.P. Kaelbling, "The Synthesis of Digital Machines with Provable Epistemic Properties," In *Proc. 1986 Conference on Theoretical Aspects of Reasoning about Knowledge,* W. Kaufmann, Los Altos, California, 1986, 83–98.

13 Nilsson, N., "Probabilistic Logic," *Artificial Intelligence,* **28** (1986) 71–88.

14 Williams, P.M., "Indeterminate Probabilities," in *Formal Methods in the Methodology of Empirical Sciences,* (M. Przelecki, K. Szaniawski and R. Wojciki, eds.), Ossolineum and D. Reidel, 1976.

15 Halpern, J.Y. and D.A. McAllester, "Likelihood, Probability, and Knowledge," *in Proc. 1984 AAAI Conference,* AAAI, Menlo Park, California, 1984, 137–141.

16 Hughes, G.E. and M.E. Creswell, *An Introduction to Modal Logic,* London: Methuen, 1968.