AN INFERENCE TECHNIQUE FOR INTEGRATING KNOWLEDGE FROM DISPARATE SOURCES

Thomas D. Garvey, John D. Lowrance, and Martin A. Fischler

SRI International, Menlo Park, California

ABSTRACT

This paper introduces a formal method for integrating knowledge derived from a variety of sources for use in "perceptual reasoning."* The formalism is based on the "evidential propositional calculus," a derivative of Shafer's mathematical theory of evidence [4]. It is more general than either a Boolean or Bayesian approach, providing for Boolean and Bayesian inferencing when the appropriate information is available. In this formalism, the likelihood of a proposition A is represented as a subinterval, [s(A),p(A)], of the unit interval, [0,1]. The evidential support for proposition A is represented by s(A), while p(A) represents its degree of plausibility; p(A) can also be interpreted as the degree to which one fails to doubt A, p(A) being equal to one minus the evidential support for "A. This paper describes how evidential information, furnished by a knowledge source in the form of a probability "mass" distribution, can be converted to this interval representation; how, through a set of inference rules for computing intervals of dependent propositions, this Information can be extrapolated from those propositions it directly bears upon, to those it indirectly bears upon; and how multiple bodies of evidential Information can be pooled* A sample application of this approach, modeling the operation of a collection of sensors (a particular type of knowledge source), illustrates these techniques.

I INTRODUCTION AND OVERVIEW

We are pursuing a program of research aimed at developing a computer-based capability for "perceptual reasoning" [2] that will make it possible to Interpret important aspects of a situation from information obtained by a collection of disparate sensors. Situational assessment Implies the need to integrate sensory information with a body of relevant "expertise," or prior knowledge. This integration poses a number of difficult technical problems that must be examined.

Among the problems focused upon in our work are the following:

* How to model sensors and other knowledge sources (KS), so as to know which situations they can provide Information about and how to interpret their responses.
* How to effectively combine (sometimes contradictory) information from multiple knowledge sources to compensate for their Individual deficiencies.
* How to automatically devise a data-acquisition/sensor-utilization strategy to maximise overall system effectiveness.

In this paper we shall concentrate on the approach to sensor modeling and knowledge integration that is currently under investigation. These form the core of the overall system.

A. Previous Work

Earlier research [1, 2, 3] led to a number of important conclusions regarding the integration of perceptual Information. First, because of the variety of knowledge types required and the particular uses of each, it became apparent that a proliferation of specialised representations was inevitable. This is a departure from standard approaches that attempt to develop a representation of sufficient scope to encompass all of the knowledge needed by a system. Use of nonmonolithic representations allows KSs to perform efficient operations on widely diverse, locally appropriate data formats. However, the problem then becomes one of somehow connecting these KSs in a flexible, effective manner.

We formulated several requirements of a reasoning paradigm for the combination and extrapolation of evidential information from disparate KSs. Whereas earlier work has focused upon a Bayesian-based probabilistic scheme, we feel that this is too restrictive. A likelihood represented by a point probability value is usually an overstatement of what is actually known, distorting the available precision.

The work described here has been jointly supported by the Defense Advanced Research Projects Agency of the Department of Defense (monitored by the Air Force Avionics Laboratory under Contract No. F33615-80-C-1110) and the Office of Naval Research under Contract No. N00014-81-C-0115.

In particular, there is no adequate, non-ad hoc representation of ignorance within a Bayeslan framework.

Another problem with a Bayeslan approach to the modeling of belief is the difficulty of ensuring and maintaining consistency in a collection of interrelated propositions. This difficulty also stems from the need to assign point probability values, even when the underlying models from which these values are derived are incapable of supplying such precise data.

There are many occasions when the inference technique of choice is probabilistic reasoning (e.g., particularly when reasoning is done with data close to the signal level), and other occasions when a (Boolean) logical formalism is preferred (e.g., when trying to combine "higher-level" knowledge). To avoid an ad hoc approach to "global" knowledge integration, the inference paradigm should flow smoothly from a probabilistic technique to a logical one, as the propositions in question become more nearly true or false. In addition, whenever the underlying model is complete and consistent enough for traditional methods to be effective, the technique should reduce to a Bayeslan paradigm.

B. The Shafer Representation

The representation we have adopted to satisfy the preceding requirements for the integration of global knowledge is based on the work of Shafer [4]. It expresses the belief in a proposition A by a sublnterval (s(A),p(A)] of the unit interval, [0,1]. The lower value, s(A), represents the "support" for that proposition and sets a minimum value for its likelihood. The upper value, p(A), denotes the "plausibility" of that proposition and establishes a maximum likelihood. Support may be Interpreted as the total positive effect a body of evidence has on a proposition, while plausibility represents the total extent to which a body of evidence falls to refute a proposition. The degree of uncertainty about the actual probability value for a proposition corresponds to the width of its interval. As will be shown, this representation with the appropriate inference rules satisfies the requirements established above.

In the remainder of this paper, we shall demonstrate Dempster's rule of combination [A] for pooling evidential information from independent knowledge sources, present an Inference mechanism for updating proposition intervals based on other dependent proposition intervals, and demonstrate their use in sensor modeling and integration.

For example, if no Information is available concerning two Initially exclusive and exhaustive possibilities, in a Bayeslan framework they are usually assigned a probability of .5. This Is quite different from specifying that nothing is known regarding such propositions.

II KNOWLEDGE REPRESENTATION AND INFERENCE

in what we call the "evidential propositional calculus," we represent a proposition using the following notation:

$$A_{[s(A),\ p(A)]},$$

where A is the proposition, s(A) the support for the proposition, and p(A) Its plausibility. p(A) is equivalent to i-s(~A), the degree to which one falls to doubt A. The Interval [s(A),p(A)I Is called the "evidential interval." The uncertainty of A, u(A), corresponds to p(A)-s(A). If u(A) is zero for all propositions, the system is Bayeslan.

The following examples illuminate some important points;

* $A_{[0,1]}$ => no knowledge at all about A.
* $A_{[0,0]}$ => A is false.
* $A_{[1,1]}$ => A is true.
* $A_{[.25,1]}$ => evidence provides partial support for A.
* $A_{[0,.85]}$ => evidence provides partial support for ~A.
* $A_{[.25,.85]}$ => probability of A is between .25 and .85; i.e., the evidence simultaneously provides support for both A and ~A.

A• Dempster's Rule of Combination

Dempster's rule is a method of integrating distinct bodies of evidence. This is most easily introduced through the familiar formalism whereby propositions are represented as subsets of a given set, here referred to as the "frame of discernment" (denoted 6). When a proposition corresponds to a subset of the frame of discernment, it is said to be "discerned." The primary advantage of this formalism is that it translates the logical notions of conjunction, disjunction, implication, and negation into the more graphic, set-theoretic notions of intersection, union, inclusion, and complementation* Dempster's rule combines evidential information expressed relative to those propositions discerned by O.

1. Single Belief Functions

We assume that a knowledge source, KS_1, distributes a unit of belief across a set of propositions for which it has direct evidence, in proportion to the weight of that evidence as it bears on each. This is represented by a function:

$$m_1: \{A_i \mid A_i \subseteq \Theta\} \dashrightarrow [0,1],$$

$$m_1(\phi) = 0,$$

$$\sum_{A_i \subseteq \Theta} m_1(A_i) = 1.$$

Those propositions are referred to as the KS's "focal elements*"

$m_1(A_i)$ represents the portion of belief that KS_1 has committed exactly to proposition A_i termed its "basic probability mass." m_1 can be depicted as a partitioned unit line segment, the length of each subsegment corresponding to the mass attributed to one of Its focal elements (Figure 1). Any mass assigned to O represents the residual "uncertainty" of the KS directly. That is, $m_1(O)$ is the mass that could not be ascribed to any smaller subset of 6 on the basis of the evidence at hand, but must Instead be assumed to be distributed in some (unknown) manner among the propositions discerned by 9. A similar Interpretation is given to mass assigned any (nonunlt) set.

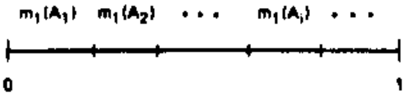

Figure 1 Probability Mass Assignment for KS_1

Once mass has been assigned to a sec of propositions, the evidential interval can be determined directly. Support for a proposition A is the total mass ascribed to A and to its subsets; the plausibility of A is one minus the sum of the mass assigned to ˜A and to subsets of "A; the uncertainty of A is equal to the mass remaining, i.e., that attributed to supersets of A, including ©.

$$s_1(A) = \sum_{A_i \subseteq A} m_1(A_i).$$

$$p_1(A) = 1 - s_1(\text{~}A).$$

$$u_1(A) = p_1(A) - s_1(A).$$

For example,

if $A = \{a\}$, $\quad A \vee B = \{a,b\}$,

$\text{~}A = \{b,c\}$, $\quad \theta = \{a,b,c\}$,

and $m_1(\langle A, \text{~}A, A \vee B, \theta\rangle) = \langle .4, .2, .3, .1\rangle$;

then $A_{[.4,\ .8]}$, $\quad A \vee B_{[.7,\ .1]}$,

$\text{~}A_{[.2,\ .6]}$, $\quad \theta_{[1,\ 1]}$.

2. Composition of Mass Functions

If the belief function of a second KS, KS_2, is also provided, the information supplied by these KSs can be pooled by computing the "orthogonal sum;" this computation is illustrated by the unit square in Figure 2.

To combine the effects of KS_1, and KS_2, we consider the unit square as representing the combined probability mass of both KSs; KS_1 partitions the square into vertical strips corresponding to its focal elements, while KS_2 partitions it into horizontal strips that correspond to its focal elements. For example, Figure 3 shows a vertical strip of measure $m_1(A_i)$ that is exactly committed to A_i by KS_1, and a horizontal strip of size $m_2(B_j)$ committed precisely to B_j by KS_2. The intersection of these strips commits exactly $m_1(A_i)m_2(B_j)$ to the combination of $A_i \cap B_j$.

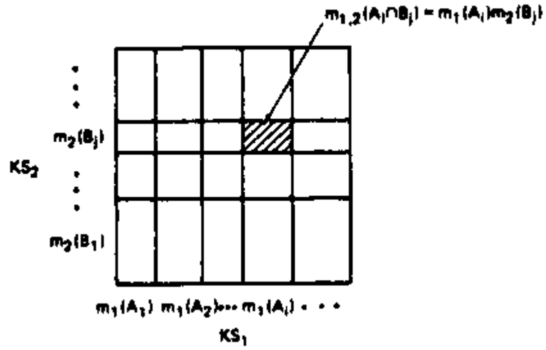

Figure 2 Composition of Mass from KS_1 and KS_2

Accordingly, we can compute the area commitment of each rectangle comprising the square. A given subset of θ, C, may have more than one rectangle exactly committed to it; the total mass allocated to C is

$$\sum_{A_i \cap B_j = C} m_1(A_i)m_2(B_j).$$

This scheme is likely to commit a portion of mass to the empty set ϕ. Every rectangle committed to $A_i \cap B_j$, where $A_i \cap B_j = \phi$, results in such a commitment. The "remedy" is to discard all such rectangles, proportionally increasing (i.e., normalizing) the size of the remaining rectangles by the following multiplicative factor:

$$N = (1 - k)^{-1},$$

where

$$k = \sum_{A_i \cap B_j = \phi} m_1(A_i)m_2(B_j),$$

thereby restoring the total probability mass to one.

There are several points of Interest with respect to Dempster's rule of combination. The operation is commutative; therefore, the order of combination is Immaterial. The operation is also associative, allowing the palrwise compositions of a sequence of KSs. When two Bayesian mass functions are combined, one associating its full unit of mass with a single proposition, the resulting support and plausibility values are the expected Bayeslan conditionals. Yet when only less precise information is available, It too can be exploited.

The degree of conflict between two KSs can be measured intuitively by the size of the factor k. The greater the value of k, the greater the degree of conflict between the two KSs. When k la one, the KSa are Irreconcilably different and the orthogonal sum doea not exist.

B. <u>Inference Rules</u>

In addition to a technique for pooling distinct bodies of evidence, rules are needed that allow evidential information to be tranalated from thoae propositions it bears upon directly to those it beara upon indirectly. These rulea are baaed on the following two principles of evidential support:

* The proposition corresponding to the frame of discernment alwaya receives full support.
* Any support committed to a proposition is thereby committed to any other proposition it implies.

From the first principle we know that s(6) - p(8) - 1. The second principle dictates that any support committed to a subset of the frame of discernment is thereby committed to Its supersets. This follows because one proposition implies another if it la a subset of that proposition in the frame of discernment. Of the total aupport committed to a given proposition A, some may be committed to one or more proper aubaeta of A, while the reat is committed exactly to A —and to no smaller subset, i.e., m(A). If it la assumed that a knowledge source expresses itself in terms of support and plausibility estimates for a selected set of propositions from the frame of discernment, a aet of inference rulea allows these estimates to be tranalated from proposition to proposition, thereby reducing uncertainty. A sampling of theae rulea follows. The statements above the line in each rule allow the statement below the line to be inferred.

$$\begin{array}{l} \\ \hline \theta[1,1]. \end{array}$$

$$\begin{array}{l} A \subset \theta \\ \hline A[0,1]. \end{array}$$

$$\begin{array}{l} A[s1(A),p1(A)] \\ A[s2(A),p2(A)] \\ \hline A[s(A),p(A)], \quad s(A) = MAX[\ s1(A),\ s2(A)\ ], \\ \qquad\qquad\qquad p(A) = MIN[\ p1(A),\ p2(A)\ ]. \end{array}$$

$$\begin{array}{l} A[s(A),p(A)] \\ \hline \sim A[s(\sim A),p(\sim A)], \quad s(\sim A) = 1 - p(A), \\ \qquad\qquad\qquad p(\sim A) = 1 - s(A). \end{array}$$

$$\begin{array}{l} A[s(A),p(A)] \\ B[s(B),p(B)] \\ \hline A \vee B[s(A\vee B),p(A\vee B)], \\ \qquad\qquad s(A\vee B) = MAX[\ s(A),\ s(B)\ ], \\ \qquad\qquad p(A\vee B) = MIN[\ 1,\ p(A) + p(B)\ ]. \end{array}$$

$$\begin{array}{l} A \vee B[s(A\vee B),p(A\vee B)] \\ A[s(A),p(A)] \\ \hline B[s(B),p(B)], \quad s(B) = MAX[\ 0,\ s(A\vee B) - p(A)\ ], \\ \qquad\qquad\qquad p(B) = p(A\vee B). \end{array}$$

$$\begin{array}{l} A[s(A),p(A)] \\ B[s(B),p(B)] \\ \hline A \& B[s(A\&B),p(A\&B)], \\ \qquad\qquad s(A\&B) = MAX[\ 0,\ s(A) + s(B) - 1\ ], \\ \qquad\qquad p(A\&B) = MIN[\ p(A),\ p(B)\ ]. \end{array}$$

$$\begin{array}{l} A \& B[s(A\&B),p(A\&B)] \\ A[s(A),p(A)] \\ \hline B[s(B),p(B)], \quad s(B) = s(A\&B), \\ \qquad\qquad\qquad p(B) = MIN[\ 1,\ 1 + p(A\&B) - s(A)\ ]. \end{array}$$

As can be easily shown, when propositions are known to be true or false (that is, when their corresponding belief Intervals become either [0,0] or [1,1]), these rules reduce to the corresponding rules of the propositional calculus. Thus, when appropriate knowledge exists, this method will enable eaay transition from a probabilistic inference computation to the standard propositional calculus.

III EXAMPLE: MODELING A KNOWLEDGE SOURCE

As our intention has been to treat sensors as specialized KSs, in this section we shell describe an approach to modeling such a KS. We shall begin by discussing the usual parameters measured by a (hypothetical) sensor, illustrating how, for this simple example, these measurements are converted to hypotheses by the inference mechanism.

A. <u>Sensor Measurements</u>

We assume thet collections of electromagnetic signal "emitters" deployed in various configurations comprise the situation of interest. Measurements of characteristics of the signals emitted by these devices will be used to formulate hypotheses about their identities. In the complete system, these hypotheses will interact with those derived from other KSs to create a more comprehensive picture of the situation. Let us first concern ourselves with a single KS and then show how it may be composed with other KSs.

1. <u>Emitter Characteristics</u>

In this example, en emitter will rediate e pulsed radar signal whose pertinent characteristics will include the carrier frequency (rf) and the pulse width (pw), which are measured directly by the receiver. For the example we assume that the emitters of interest ere of types El, E2, E3, B4, or B5. The goal of the program is to identify a signal as having originated from one of those types.

The Information about the parameter values likely to be exhibited by an emitter is presented in the form of parameter distribution graphs—for example, as shown in Figure 3. These curves indicate the probability that any given emitter (of the type indicated) will have a specific parameter value; the total area under each curve is one.

A typical approach to identifying an emitter is to look up the measured parameters in a table. In addition to difficulties traceable to the static nature of the table (e.g., emitter characteristics are not expected to remain stable and constant in actual operation), the technique gives little Information regarding the relative likelihoods of ambiguous identifications.

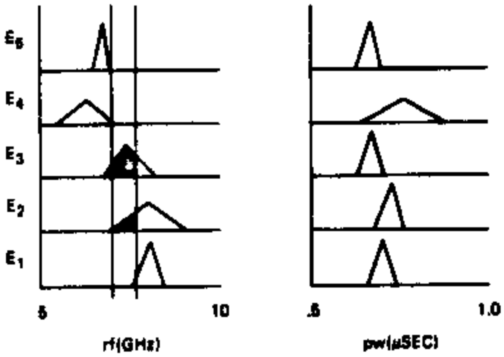

Figure 3 Examples of Typical Emitter Parameter Distributions

2. Sensor Characteristics

A sensor (receiver) will specify a range of possible values for the measured emitter parameter, as determined by the resolution of the sensor's measurements. For example, the receiver may specify an emitter's frequency value as lying within a band of frequencies, say from 5.5 to 5.6 GHz. Similarly, other emitter parameters will also be specified as falling within a range of values.

Our previous work [1] described how sensor models were modified in the event of changing environmental conditions. In the approach described here, such environmental factors will instead determine the total mass a sensor may allocate to propositions other than 3. In effect, the uncertainty U of a receiver in the prevailing conditions is its minimal commitment to 6 (I.e., m(O) > U), leaving only (1 - U) of the mass to be freely distributed.

B. Modeling the Operation of a Sensor

The modeling process begins with the determination of a frame of discernment. If the task Is to determine the true value of some variable, the frame of discernment is the set of all possible values for that variable. For the problem at hand, each element of the frame of discernment consists of an emitter type paired with a feature vector representing one possible electromagnetic signature that such an emitter might exhibit. Thus, Θ is a subset of all of the combinations of emitter types (ET), radio frequencies (RF), and pulse widths (PW).

$$\Theta \subseteq ET \times RF \times PW.$$

The key requirement of the frame of discernment is that all the propositions of interest be in correspondence with Its subsets. In the current context the following propositions are some of those that might be of interest:

(The emitter is type E1)
$= \{q \mid q \in \Theta \text{ and } et(q) = E1\};$

(The radio frequency is between 5.5 and 5.6 GHz)
$= \{q \mid q \in \Theta \text{ and } 5.5 \leq rf(q) \leq 5.6\};$

(The emitter is type E1 with pulse width between .68 and .7 μs)
$= \{q \mid q \in \Theta \text{ and } et(q) = E1 \text{ and } .68 \leq pw(q) \leq .7\}.$

Once a frame of discernment has been determined, it can be represented as a dependency graph [5]. In this formalism propositions are represented by nodes, their interrelationships by arcs. These interrelationships can be Interpreted either as set-theoretic notions relative to the frame of discernment (e.g., Intersection, union, inclusion, and complementation), or as logical connectives (e.g., conjunction, disjunction, implication, and negation). The appropriate subset of propositions and relationships, so represented, depends on the preferred vocabulary of discourse among the KSs. Those propositions to which the KSs tend to assign mass need to be included, along with those relationships that best describe their interdependence. Once such a dependency graph has been established, it provides an integrated framework for both the combination and extrapolation of evidential Information.

In the current context there is a subgraph for each emitter feature. At the lowest level of these subgraphs is a set of propositions representing the smallest bands into which that continuous feature has been partitioned—this partitioning being necessary within a propositional framework. These primitive bands form the basis of a hierarchy in each subgraph, relating larger bands to more primitive ones. The emitter types are similarly represented, the higher elements in the hierarchy corresponding to disjunctions of emitter types. All this is tied together by one last subgraph that relates the base elements of the hierarchies to elements of the frame of discernment, the frame of discernment consisting of the possible combinations of base elements. Figure 4 is a sketch of this dependency graph, with each node representing a proposition equal to the disjunction of those immediately below it, and the conjunction of those immediately above it.

This dependency graph contains all the information needed to determine the collective impact of several bodies of evidence on all the

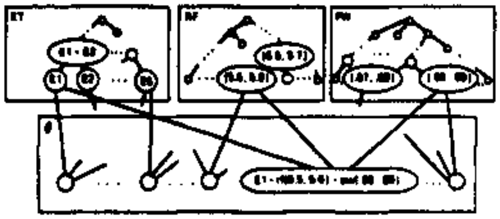

Figure 4 Dependency Graph for Sensor Model

propositions of interest. Given several distinct bodies of evidential Information extracted from the environment by several knowledge sources, repeated applications of Dempster's rule followed by repeated applications of the inference rules for support and plausibility propagation—all based on the information embodied in the dependency graph—results in a support and plausibility estimate for every proposition of Interest [5]. There are no restrictions regarding which propositions serve as premises or conclusions. Information about radio frequency and pulse width can be used to determine the most likely types of emitters—or information about emitters and pulse width can be used to predict the expected radio frequency. Inferencing is unconstrained.

C. Simplification of the Sensor Model

In the preceding discussions, we showed how inferences could be drawn in a formal system that modeled all relevant elements of 6. It is frequently inconvenient to model these elements individually. For example, too many elements may be needed to represent the resolution of any particular sensor. An obvious simplification is to compute new propositions in 9 as they are needed, ror example, when a receiver reports a signal in a specific frequency band, propositions can then be created which assert that the signal originated from one element of a subset of possible emitter types. The exact hypotheses and their associated mass allocations are determined by comparing receiver measurements with tabulated information about the emitters.

1. Initial Mass Computations

The first step is to convert sensory measurements into a probability mass distribution over propositions. In essence, the parameter measurement range is overlaid on the curves representing distributions of emitter parameters, as shown in Figure 3. The area of the distribution curve is computed for each emitter (propositions are created only for those emitters whose parameter ranges overlap the sensor's report). A set of "basic mass numbers" is then computed by normalising the resultant areas to bring their total area to one. This process is exactly equivalent to computing the probability of each emitter, conditioned upon the measured parameter's falling in the specified range (and assuming that only the tabulated emitter could radiate the received signal).

The uncertainty U of the receiver is accounted for through reducing each basic mass number by multiplying it by a factor equal to one minus U. This new set of mass numbers then represents the contribution of the receiver measurement to the support of the proposition.

2. Example

In this example we assume that there are five emitter types {El, ..., E5}, whose rf and pw characteristics are shown graphically in Figure 3. The receiver has reported a frequency measurement of 7.6 to 7.7 GHz and a pulse width range of •68 to .7 μs. Assuming an uncertainty of .3 in the rf measurement and an uncertainty of .2 for pw, the resulting mass functions are

$$m_{rf}(\langle E1,E2,E3,E4,E5\rangle) = \langle .13,.22,.35,0,0\rangle$$

and

$$m_{pw}(\langle E1,E2,E3,E4,E5\rangle) = \langle .26,.085,.17,.034,.26\rangle.$$

Combining these with Dempster's rule gives the composite mass function,

$$m_{rf\&pw}(\langle E1,E2,E3,E4,E5\rangle) = \langle .25,.16,.33,.018,.14\rangle,$$

with a resulting uncertainty of 0.11. This computation is illustrated in Figure 5, in which all rectangles attributed to ϕ are shaded and the remaining rectangles labeled with the proposition receiving that mass. These values convert directly to intervals on the propositions, as shown:

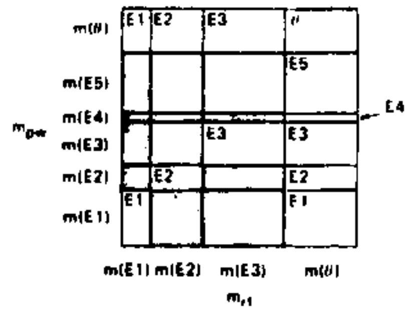

Figure 5 Composition of the Mass Assignments by Sensor Models

E1[.25, .36]
E2[.16, .27]
E3[.33, .43]
E4[.018, .13]
E5[.14, .25]

This information may be readily combined with information provided by other KSs. For example, a KS that indicated a high likelihood of encountering an El might produce the following mass function:

$$m_{prior}(\langle E1,E2,E3,E4,E5\rangle) = \langle .7,0,0,0,0\rangle,$$

which, when integrated with the receiver measurements, would result in a mass function

$$m_{composite}(\langle E1,E2,E3,E4,E5\rangle) = \langle .59,.089,.18,.01,.076\rangle,$$

with an uncertainty of .06. This leads to the following relevant hypotheses:

$$E1[.59,.65] \text{ and } E3[.18,.24].$$

Based on sensor data alone, the method leads to two primary hypotheses, Elr 25 ,351 and E3r 33 43]. E3 is slightly favored over El. When external evidence is brought to bear, the support for E1 becomes significantly greater than for all others (and, in fact, all others except E3 drop to very low levels of support). Any other KS that provides a mass assignment over this set of propositions may also be combined.

This simplification of the formal method provides the ability to integrate information quickly from a variety of sources, even in those areas where the necessary propositions have not already been extracted from 6. The technique does not yet allow the propagation of evidence to arbitrarily selected propositions from the network. For example, it is not possible to take the structure defined for this problem and use It to determine what radio frequency values should be expected on the basis of pulse width data—a process easily carried out by the full representation. This is an area of current research. A related computational technique, restricted to evidence that either confirms or denies a single proposition, is also being investigated [6].

IV SUMMARY

We have briefly described an inference technique that appears to satisfy many of the requirements for reasoning in perceptual domains. In particular, the method provides the capability for (Bayeslan) probabilistic reasoning when the appropriate underlying models are available (e.g., at the lowest levels of the system), (evidential) subjective reasoning when incomplete descriptions must be used (e.g., at the "middle" levels of the system), and (Boolean) logical reasoning when the truth values of propositions are true and false. This technique allows us to augment a static, Incomplete model with current sensory information.

The approach provides a formal technique for updating the likelihoods of propositions in a consistent manner. In effect, by simultaneously performing computations over a collection of propositions, the method maintains global consistency without the problems frequently plaguing techniques that perform iterative updating by means of local rules. Most importantly, besides offering an inference technique that can be used within a KS (as illustrated), the method provides a "language" for KSa to communicate with one another, as well as furnishing the means for linking disparate sources of information.

In our previous work on perceptual-reasoning systems, we evolved a number of effective generic (e.g., terrain, weather, etc.) and domain-specific (e.g., sensor) KSs. Our current research, focusing on the evidential proposltlonal calculus as the integrating medium, alms at developing a general framework for linking these KSs together smoothly and flexibly.

REFERENCES

1. T. D. Garvey and M. A. Fischler, "Machine-Intelligence-Based Multisensor ESM System," Technical Report AFAL-TR-79-1162, Air Force Wright Avionics Laboratory, Wright-Patterson Air Force Base, Ohio (October 1979).

2. T. D. Garvey and M. A. Fischler, "Perceptual Reasoning in a Hostile Environment," Proceedings of the First Annual Conference on Artificial Intelligence, Stanford University, Stanford, California, pp. 253-255 (August 1980).

3. T. D. Garvey and M. A. Fischler, "The Integration of Multi-Sensor Data for Threat Assessment," Proceedings of the Fifth Joint Conference on Pattern Recognition, Miami Beach, Florida, pp. 343-347 (December 1980).

4. G. Shafer, A Mathematical Theory of Evidence (Princeton University Press, Princeton, New Jersey 1976).

5. J. D. Lowrance, "Dependency-Graph Models of Evidential Support," Ph.D. Dissertation, University of Massachusetts, Amherst, Massachusetts (in preparation).

6. B. A. Barnett, "Computational Methods for a Mathematical Theory of Evidence," Proceedings of the Seventh International Joint Conference on Artificial Intelligence, Vancouver, British Columbia, Canada (1981).