A FRAMEWORK FOR EVIDENTIAL-REASONING SYSTEMS*

John D. Lowrance Thomas D. Garvey
Thomas M. Strat

Artificial Intelligence Center
SRI International
Menlo Park, California

Abstract

Evidential reasoning is a body of techniques that supports automated reasoning from evidence. It is based upon the Dempster-Shafer theory of belief functions. Both the formal basis and a framework for the implementation of automated reasoning systems based upon these techniques are presented. The formal and practical approaches are divided into four parts (1) specifying a set of distinct propositional spaces, each of which delimits a set of possible world situations (2) specifying the interrelationships among these propositional spaces (3) representing bodies of evidence as belief distributions over these propositional spaces and (4) establishing paths for the bodies of evidence to move through these propositional spaces by means of evidential operations, eventually converging on spaces where the target questions can be answered.

1 Introduction

For the past several years, we have been addressing perceptual problems that bridge the gap between low-level sensing and high-level reasoning [Low82,GLF81,LG83b,LG83a,LSG86,Wes86]. Problems that fall into this gap are often characterized by multiple evidential sources of real-time data, which must be properly integrated with general knowledge about the world to provide an understanding of the situation that is sufficiently rich to support high-level goals. In this paper, we describe a formal framework for reasoning with perceptual data that forms the basis for evidential-reasoning[1] systems.

The information required to understand the current state of the world comes from multiple sources: real-time sensor data, previously stored general knowledge, and current contextual information. Sensors typically provide *evidence* in support of certain conclusions. Evidence is characteristically uncertain: it allows for multiple possible explanations; it is incomplete: the source rarely has a full view of the situation; and it may be completely or partially incorrect. The quality and the ease with which situational information may be extracted from a synthesis of current sensor data and prestored knowledge is a function both of how strongly the characteristics of the sensed data focus on appropriate intermediate conclusions and on the strength and effectiveness of the relations between those conclusions and situation events.

Given its characteristics, evidence is not readily represented either by logical formalisms or by classical probabilistic estimates. Because of this, developers of automated systems that must reason from evidence have frequently turned to informal, heuristic methods for handling uncertain information. The "probabilities" produced by these informal approaches often cause difficulties in interpretation. The lack of a formally consistent method can cause problems in extending the capabilities of such systems effectively. Our work in evidential reasoning was motivated by these shortcomings. Our theory is based on the Shafer-Dempster theory of evidence [Dem68,Sha76,Sha86] and aims to overcome some of the difficulties in reasoning from evidence by providing a natural representation for evidential information, a formal basis for drawing conclusions from evidence, and a representation for belief.

In evidential reasoning, a *knowledge source* (KS) is allowed to express probabalistic opinions about the (partial) truth or falsity of statements composed of subsets of propositions from a space of distinct, exhaustive possibilities (called the *frame of discernment*). The theory allows a KS to assign belief to the individual propositions in the space or to disjunctions of these propositions or both. When it assigns belief to a disjunction, a KS is explicitly stating that it does not have enough information to distribute this belief more precisely. This condition has the attractive feature of enabling a KS to distribute its belief to statements whose granularity is appropriate to its state of knowledge. Also, the statements to which belief is assigned are not required to be distinct from one another. The distribution of beliefs over a frame of discernment is called a *body of evidence*.

Evidential reasoning provides a formal method, *Dempster's Rule of Combination*, for fusing (i.e., pooling) two bodies of evidence. The result is a new body of evidence representing the consensus of the two original bodies of evidence, which may in turn be combined with other evidence. Because belief may be associated directly with a disjunction of propositions, the probability in any selected proposition is typically underconstrained. This necessitates an interval measure of belief, because belief associated with a disjunction may, based upon additional information, devolve entirely upon any one of the disjuncts. Thus, an interval associated with a proposition implies that the true probability associated with that proposition must fall somewhere in the interval. A side-effect of applying Dempster's rule is a measure of *conflict* between the two bodies of evidence that provides a means for detecting possible gross errors in the information.

*This research was sponsored in part by the U.S. Navy Space and Naval Warfare Systems Command and the Defense Advanced Research Project Agency under contract N00039-83-K-0656 and by the U.S. Army Signal Warfare Center under contract DAAL02-85-C-0082.

[1] *Evidential reasoning* is a term coined by SRI International [LG82] to denote the body of techniques specifically designed for manipulating and reasoning from evidential information as characterized in this paper.

Current expert-systems technology is most effective when domain knowledge can be modeled as a set of loosely interconnected concepts (i.e., propositions) [DK77]; this loose interconnection justifies an *incremental* approach to updating beliefs. In most of our work, there is the potential for strong interconnectivity among beliefs in propositions. We, therefore, focus on a body of evidence as a primitive, meaningful collection of interrelated (dependent) beliefs; updating the belief in one proposition affects the entire body of evidence (other work has addressed the concept of a body of evidence in a production-rule formalism [Kon79,LB82] by creating special entities).

Evidential reasoning provides options for the representation of information: independent opinions are expressed by multiple (independent) bodies of evidence; dependent opinions (in which belief in one proposition depends on that of another) can either be expressed by a single body of evidence or by a network that describes the interrelationships among several bodies of evidence. These networks of bodies of evidence capture the geneology of each body (similar in spirit to those of [Coh85]) and are used in a manner similar to data-flow models [WA84] updating interrelated beliefs (i.e., for belief revision [Doy81]).

In this paper we assume some familiarity with the Dempster-Shafer theory of beliefs, although the appropriate equations from this theory are included. We begin with a discussion of the formal approach to the problem of reasoning from evidence and then progress to a description of the implementation approach, including an example. We close with a short description of the system that we have developed for applying evidential reasoning.

2 Formal Approach

2.1 Framing the Problem

The first step in applying evidential reasoning to a given problem is to delimit a propositional space of possible situations. Within the theory of belief functions, this propositional space is called the *frame of discernment*. It is so named because all bodies of evidence are expressed relative to this surrounding framework, and it is through this framework that the interaction of the evidence is discerned. A frame of discernment delimits a set of possible situations, exactly one of which is true at any one time. For example, the problem to be addressed is that of locating a ship. In this case, the frame of discernment consists of the set of all possible locations for that vessel. This might be represented by a set Θ_A in which each element a_i corresponds to a possible location:

$$\Theta_A = \{a_1, a_2, \ldots, a_n\} \quad .$$

Once a frame of discernment has been established, propositional statements can be represented by disjunctions of elements from the frame corresponding to those situations for which the statements are true. For example, the proposition A_i might correspond to the statement that the vessel is located in port, in which case A_i would be represented by the subset of elements from Θ_A that correspond to possible locations within port facilities:

$$\mathsf{A}_i \subseteq \Theta_A \quad .$$

Other propositions related to locating this vessel can be similarly represented as subsets of Θ_A (i.e., as elements of the power set of Θ_A, denoted 2^{Θ_A}). Once this has been accomplished, logical questions can be posed and resolved in terms of the frame. Given two propositions, A_i and A_j, the following logical operations and relation can be resolved through the associated set operations and relation:

$$\begin{aligned} \neg \mathsf{A}_i &\iff \Theta_A - \mathsf{A}_i \\ \mathsf{A}_i \wedge \mathsf{A}_j &\iff \mathsf{A}_i \cap \mathsf{A}_j \\ \mathsf{A}_i \vee \mathsf{A}_j &\iff \mathsf{A}_i \cup \mathsf{A}_j \\ \mathsf{A}_i \Rightarrow \mathsf{A}_j &\iff \mathsf{A}_i \subseteq \mathsf{A}_j \quad . \end{aligned}$$

If other aspects of ships are of interest besides their location, then additional frames of discernment might be defined. For example, the activities of these ships might be of interest. If so, an additional frame Θ_B might be defined to include elements corresponding to refueling, loading cargo, unloading cargo, being enroute, and the like. Propositional statements pertaining to a ship's activity can then be defined relative to this frame; e.g.,

$$\begin{aligned} \Theta_B &= \{b_1, b_2, \ldots, b_n\} \\ \mathsf{B}_j &\subseteq \Theta_B \quad . \end{aligned}$$

So far, propositional statements pertaining to a ship's location or pertaining to its activity can be addressed separately, but they cannot be jointly considered. To do this, one must first define a *compatibility relation* between the two frames. A compatibility relation simply describes which elements from the two frames can be true simultaneously. For example, a ship located at a loading dock might be loading or unloading cargo, but is not refueling, or enroute. In other words, being located at a loading dock is only compatible with one of two activities, loading or unloading. Thus, the compatibility relation between frames Θ_A and Θ_B is a subset of the cross product of the two frames. A pair (a_i, b_j) is included if and only if they can be true simultaneously. There is at least one pair (a_i, b_j) included for each a_i in Θ_A (the analogue is true for each b_j):

$$\Theta_{A,B} \subseteq \Theta_A \times \Theta_B \quad .$$

Using the compatibility relation $\Theta_{A,B}$ we can define a *compatibility mapping* $C_{A \mapsto B}$ for translating propositional statements expressed relative to Θ_A to statements relative to Θ_B. If a statement A_k is true, then the statement $C_{A \mapsto B}(\mathsf{A}_k)$ is also true:

$$\begin{aligned} C_{A \mapsto B} : 2^{\Theta_A} &\mapsto 2^{\Theta_B} \\ C_{A \mapsto B}(\mathsf{A}_k) &= \{b_j | (a_i, b_j) \in \Theta_{A,B}, a_i \in \mathsf{A}_k\} \quad . \end{aligned}$$

Instead of translating propositional statements between these two frames via $C_{A \mapsto B}$ and $C_{B \mapsto A}$, we might choose to translate these statements to a common frame that captures all of the information. This common frame is identical to the compatibility relation $\Theta_{A,B}$. Frame Θ_A (and analogously Θ_B) is trivially related to frame $\Theta_{A,B}$ via the following compatibility relation and compatibility mappings:

$$
\begin{aligned}
\Theta_{A,(A,B)} &= \{(a_i,(a_i,b_j))|(a_i,b_j) \in \Theta_{A,B}\} \\
C_{A\mapsto(A,B)}(\mathsf{A}_k) &= \{(a_i,b_j)|(a_i,(a_i,b_j)) \in \Theta_{A,(A,B)}, a_i \in \mathsf{A}_k\} \\
&= \{(a_i,b_j)|(a_i,b_j) \in \Theta_{A,B}, a_i \in \mathsf{A}_k\} \\
C_{(A,B)\mapsto A}(\mathsf{X}_k) &= \{a_i|(a_i,b_j) \in \Theta_{A,B}, (a_i,b_j) \in \mathsf{X}_k\} \quad .
\end{aligned}
$$

Clearly, as more aspects of these ships become of interest, the number and complexity of the frames and compatibility mappings increases. However, there is a trade-off between the complexity of individual frames and the complexity of the network of compatibility mappings connecting them. We might define a single (complex) frame that encompasses all aspects of interest or, alternatively, define a (complex) network of frames that includes a distinct frame for each aspect of interest. Of course, these may not be equivalent. For example, consider the following frame:

$$\Theta_{A,B,C} = \{(a_1,b_1,c_1),(a_2,b_1,c_2),(a_2,b_2,c_2)\} \quad .$$

If this frame properly captures the relationship among frames Θ_A, Θ_B, and Θ_C, then c_1 is the only element from Θ_C compatible with a_1 from Θ_A. However, if we maintain these as three separate frames connected by compatibility mappings, $C_{A\mapsto B}, C_{B\mapsto A}, C_{B\mapsto C}$, and $C_{C\mapsto B}$, both c_1 and c_2 are compatible with a_1 because a_1 is compatible with b_1, and b_1 is compatible with both c_1 and c_2; i.e., $C_{B\mapsto C}(C_{A\mapsto B}(\{a_1\})) = \{c_1, c_2\}$. However, if a_1 is true, then it follows that either c_1 or c_2 is true. Thus, the reasoning based on a well-formed *gallery* of interconnected frames is sound but not necessarily complete. A gallery is well formed if there exists a single all encompassing frame whose answers are always included in the answers based upon the gallery.

In dynamic environments, compatibility relations can be used to reason over time. If Θ_{A1} represents the possible states of the world at time one and Θ_{A2} represents the possible states at time two, then a compatibility relation, $\Theta_{A1,A2}$, can capture the possible state transitions. For example, Θ_{A1} and Θ_{A2} might both represent the possible locations of a ship (i.e., they are identical to Θ_A as previously defined), then $\Theta_{A1,A2}$ could represent the constraints on that ship's movement. A pair of locations (a_i, a_j) would be included in $\Theta_{A1,A2}$ if a ship located at a_i on Day 1 (i.e., time) could reach a_j by Day 2. If we assume that the possible movements of a ship are constrained in the same way over any two day period, then the compatibility mapping associated with this compatibility relation can be reapplied as many times as necessary to constrain the possible locations of a ship across an arbitrary number of days.

2.2 Analyzing the Evidence

Once a gallery has been established, the available evidence can be analyzed. The goal of this analysis is to establish a line of reasoning, based upon both the possibilistic information in the gallery and the probabilistic information from the evidence that determines the most likely answers to some questions. The gallery delimits the space of possible situations, and the evidential information establishes the likelihoods of these possibilities. Within an analysis, bodies of evidence are expressed relative to frames in the gallery, and paths are established for the bodies of evidence to move through the frames via the compatibility mappings. An analysis also specifies if other evidential operations are to be performed, including whether multiple bodies of evidence are to be combined when they arrive at common frames. Finally, an analysis specifies which frame and ultimate bodies of evidence are to be used to answer each target question. Thus, an analysis specifies a means of arguing from multiple bodies of evidence towards a particular (probabilistic) conclusion. An analysis, in an evidential context, is the analogue of a proof tree in a logical context.

To begin, each body of evidence is expressed relative to a frame in the gallery. Each is represented as a mass distribution (e.g., m_A) over propositional statements discerned by a frame (e.g., Θ_A):

$$
\begin{aligned}
m_A : 2^{\Theta_A} &\mapsto [0,1] \\
\sum_{\mathsf{A}_i \subseteq \Theta_A} m_A(\mathsf{A}_i) &= 1 \\
m_A(\emptyset) &= 0 \quad .
\end{aligned}
$$

Intuitively, mass is attributed to the most precise propositions a body of evidence supports. If a portion of mass is attributed to a proposition A_i, it represents a minimal commitment to that proposition and all the propositions it implies. Additional mass attributed to a proposition A_j that is compatible with A_i, but does not imply it (i.e., $\emptyset \neq \mathsf{A}_i \cap \mathsf{A}_j \neq \mathsf{A}_j$), represents a potential commitment: mass that neither supports nor denies that proposition at present but might later move either way based upon additional information.

To *interpret* this body of evidence relative to the question A_j, we calculate its *support* and *plausibility* to derive its *evidential interval* as follows:

$$
\begin{aligned}
Spt(\mathsf{A}_j) &= \sum_{\mathsf{A}_i \subseteq \mathsf{A}_j} m_A(\mathsf{A}_i) \\
Pls(\mathsf{A}_j) &= 1 - Spt(\Theta_A - \mathsf{A}_j) \\
[Spt(\mathsf{A}_j), Pls(\mathsf{A}_j)] &\subseteq [0,1] \quad .
\end{aligned}
$$

The lower bound of an evidential interval indicates the degree to which the evidence supports the proposition, while the upper bound indicates the degree to which the evidence fails to refute the proposition, i.e., the degree to which it remains plausible. This evidential interval, for the most part, corresponds to bounds on the probability of A_j. Thus, complete ignorance is represented by an evidential interval of $[0.0, 1.0]$ and a precise probability assignment is represented by the "interval" collapsed about that point (e.g., $[0.7, 0.7]$). Other degrees of ignorance are captured by evidential intervals with widths other than 0 or 1 (e.g., $[0.6, 0.8], [0.0, 0.5], [0.9, 1.0]$).

If a body of evidence is to be interpreted relative to a question expressed over a different frame from the one over which the evidence is expressed, a path of compatibility relations connecting the two frames is required. The mass distribution expressing the body of evidence is then repeatedly *translated* from frame to frame, via compatibility mappings, until it reaches the ultimate frame of the question. In translating m_A from frame Θ_A to frame Θ_B via compatibility mapping $C_{A\mapsto B}$, the following computation is applied to derive the translated mass distribution m_B:

$$m_B(\mathsf{B}_j) = \sum_{C_{A\mapsto B}(\mathsf{A}_i)=\mathsf{B}_j} m_A(\mathsf{A}_i)$$

Intuitively, if we (partially) believe A_i, and A_i implies B_j, then we should have the same (partial) belief in B_j. This same method is applied to move mass distributions among frames that represent states of the world at different times. However, when this is the case, the operation is called *projection.*

Once two mass distributions m_A^1 and m_A^2 representing independent opinions are expressed relative to the same frame of discernment, they can be *fused* (i.e., combined) using *Dempster's Rule of Combination.* Dempster's rule pools mass distributions to produce a new mass distribution m_A^3 that represents the consensus of the original disparate opinions. That is, Dempster's rule produces a new mass distribution that leans towards points of agreement between the original opinions and away from points of disagreement. Dempster's rule is defined as follows:

$$\begin{aligned} m_A^3(\mathsf{A}_k) &= (1-k)^{-1} \sum_{\mathsf{A}_i\cap \mathsf{A}_j=\mathsf{A}_k} m_A^1(\mathsf{A}_i)m_A^2(\mathsf{A}_j) \\ k &= \sum_{\mathsf{A}_i\cap \mathsf{A}_j=\emptyset} m_A^1(\mathsf{A}_i)m_A^2(\mathsf{A}_j) \neq 1 \quad . \end{aligned}$$

Since Dempster's rule is both commutative and associative, multiple (independent) bodies of evidence can be combined in any order without affecting the result. If the initial bodies of evidence are independent, then the derivative bodies of evidence are independent as long as they share no common ancestors. Thus, in the course of constructing an analysis, attention must be paid to the way that evidence is propagated and combined to guarantee the independence of the evidence at each combination.

Other evidential operations can also be included in an analysis. One frequently used operation is *discounting.* This operation adjusts a mass distribution to reflect its source's credibility (expressed as a discount rate $r \in [0,1]$). If a source is completely reliable ($r = 0$), discounting has no effect; if it is completely unreliable ($r = 1$), discounting strips away all apparent information content; otherwise, discounting lowers the apparent information content in proportion to the source's unreliability:

$$m_A^{\%}(\mathsf{A}_i) = \begin{cases} (1-r)m_A(\mathsf{A}_i), & \mathsf{A}_i \neq \Theta_A \\ r + (1-r)m_A(\Theta_A), & \text{otherwise} \end{cases} .$$

Other evidential operations include *summarization* and *gisting* (among others). Summarization eliminates extraneous details from a mass distribution by collecting all of the extremely small amounts of mass attributed to propositions and attributing the sum to the disjunction of those propositions. Gisting produces the "central" Boolean-valued statement that captures the essence of a mass distribution. This is particularly useful when explaining lines of reasoning.

3 Implementation Approach

In implementing this formal approach, we have found that the gallery, frames, compatibility relations, and analyses can all be represented straightforwardly as graphs consisting of nodes connected by directed edges. This has led us to use **Grasper II**^{TM} [Low86,Low78], a programming language extension to LISP that introduces graphs as a primitive data type. A graph in Grasper II consists of a set of labeled subgraphs. Each subgraph consists of a set of labeled nodes and a set of labeled, directed edges that connect pairs of nodes. Each node, edge, and subgraph have values that can be used as general repositories for information. Once the graphical representations have been established for the gallery, frames, compatibility relations, and analyses, the remainder of the formal approach is easily implemented.

The first step is to define the gallery. If the problem is to reason about the locations and activities of ships, we might include two frames: a LOCATIONS frame and an ACTIVITIES frame. These are each represented as nodes in a subgraph called the SHIP-GALLERY (Figure 1). In addition, the gallery might include three compatibility relations represented as edges. One compatibility relation, LOCATIONS-ACTIVITIES, relates locations to activities and is represented by an edge from LOCATIONS to ACTIVITIES. The two other compatibility relations, DELTA-LOCATIONS and DELTA-ACTIVITIES, describe how a ship's location and activity on one day are related to the next day's. Each of these is represented by an edge that begins and ends at the same node.

The next step is to define the frames in the gallery. Each of these is represented by a subgraph sharing the same name as a node from the gallery. Each such subgraph includes a node for each element of the frame and may include additional nodes representing aliases, i.e., named disjunctions of elements. Each of these additional nodes have edges pointing to elements of the frame (or other aliases) that make up the disjunction. The LOCATIONS frame (Figure 2) includes six elements (ZONE1, ZONE2, ZONE3, CHANNEL, LOADING-DOCK, REFUELING-DOCK) and three aliases (IN-PORT, DOCKED, AT-SEA). The ACTIVITIES frame (Figure 3) includes five elements (ENROUTE, TUG-ESCORT, UNLOADING, LOADING, REFUELING).

Each compatibility relation in the gallery is represented as a subgraph that includes the nodes from the frames that they relate with edges connecting compatible elements. For example, in the LOCATIONS-ACTIVITIES compatibility relation (Figure 4), ZONE1, ZONE2, and ZONE3 are all connected to ENROUTE (becuase these zones represent areas at sea), CHANNEL is connected to TUG-ESCORT (because a ship entering or leaving the port at the end of this channel would be under tugboat control), LOADING-DOCK is connected to both LOADING and UNLOADING (because either activity is consistent with being at that dock), and REFUELING-DOCK is connected to REFUELING. DELTA-LOCATIONS and DELTA-ACTIVITIES (Figures 5 and 6) relate frames to themselves. They represent possible state transitions in their respective frames over any two day period. Edges connect compatible elements from one day to the next. DELTA-LOCATIONS indicates that the zones are linearly ordered and that a ship must pass through the channel to get to either the loading or refueling docks. It also indicates that a ship will only remain at the refueling dock or in the channel for one day at a time but may remain anywhere else for any number of days. In DELTA-ACTIVIES it can be seen that a ship must progress through TUG-ESCORT from ENROUTE before proceeding to REFUELING or UNLOADING and that REFUELING and TUG-ESCORT are one day activities. Further, a ship must go through LOADING after UNLOADING before returning to TUG-ESCORT.

After the gallery and its supporting frames and compatibility relations have been established, evidential analyses can be constructed. These analyses are represented as data-flow graphs where the data and the operations are evidential. Figure 7 is one such analysis. Here primitive bodies of evidence are represented by elliptical nodes and derivative bodies of evidence are represented by circular nodes. Diamond-shaped nodes represent interpretations of bodies of evidence. The values of these nodes are used as repositories for the information (i.e., data) that they represent (Figure 8). For bodies of evidence this includes a frame of discernment (including the day to which the evidence pertains), a mass distribution, and other supporting information. Edges pointing to a derivative node are labeled with the evidential operation that is applied to the bodies of evidence, at the other ends of the edges, to derive the body of evidence represented by this node.

In the analysis of a ship in Figure 7, there are three primitive bodies of evidence. REPORT1 locates the ship on Day 1 saying that there is a 70 percent chance that it can be found in the CHANNEL and a 30 percent chance that it is in ZONE1; REPORT2 says that the ship was IN-PORT on Day 2; and REPORT3 indicates that the ship was LOADING cargo on Day 3. REPORT1 is taken at face value, but REPORT2 and REPORT3 have been discounted by 20 percent and 40 percent, respectively, to derive D2 and D3, reflecting doubt in the credibility of these reports. REPORT1 has been projected forward by one day to derive P1 [2] and then has been fused with D2 to derive a consensus for Day 2, F12. D3 has been projected backwards in time by one day to derive P3 and then has been translated from the ACTIVITIES frame to the LOCATIONS frame. Finally, this result, T3, has been fused with F12 to derive a consensus, based on all three reports, about the ship's location on Day 2.

The interpretation nodes in this analysis track the evidential intervals for some key propositions. I1 is based soley on REPORT1 and indicates that there is precisely a 70 percent chance of the ship being IN-PORT[0.7, 0.7] and no chance of it being DOCKED [0.0,0.0]on Day 1. IP1 indicates that, based soley upon REPORT1, after one day has ellapsed, nothing is known about whether the ship is IN-PORT [0.0, 1.0], but that it may now be DOCKED [0.0,.7.0]. If REPORT2 is included after being discounted, IF12 indicates that there is strong reason to believe that the ship is IN-PORT [0.8, 1.0], but there is conflicting information concerning whether or not it is DOCKED [0.56,0.7]. IT3 indicates that based soley upon REPORT3, after having been discounted, projected backwards a day, and translated to the LOCATION frame, that there is 0.6 support and 1.0 plausibility for both IN-PORT and DOCKED. Finally, when all three reports are considered, IF123 indicates strong belief that the ship is IN-PORT [0.9, 1.0] on Day 2 and a reasonably strong belief, though mixed, that it is also DOCKED [0.78,0.85].

[2] Note that the distribution at REPORT1 is a Bayesian distribution (i.e., a distribution over exclusive elements), but application of the projection operation results in a non-Bayesian distribution at P1.

4 Evidential-Reasoning Systems

To support the construction, modification, and interrogation of evidential analyses, we have developed **Gister**^{TM}. Gister supports an interactive, menu-driven, graphical interface that allows these structures to be easily manipulated. The user simply selects from a menu to add an evidential operation to an analysis, to modify operation parameters (e.g., discount rates), or to change any portion of a gallery including its frames and compatibility relations. In response, Gister updates the analyses.

All of the figures in this paper are actual screen images from Gister. Figure 7 includes the menus for working with analyses. On the left side of the screen is a menu of nouns. The user determines with what class of objects he wishes to work and selects the appropriate noun from the menu. Once a noun has been selected, a menu of verbs appears on the right side of the screen. A selection from this menu invokes the operation corresponding to the selected verb on the previously selected noun. The user then designates the appropriate nodes, edges, and the like for the selected operation.

Unlike other expert systems, Gister is designed as a tool for the domain expert. With this tool, an expert can quickly and flexibly develop a line of reasoning specific to a given domain situation. This differs markedly from other expert systems in which a single line of reasoning is developed by an expert and then is instantiated over different situations by nonexperts.

This approach has been successfully applied to Naval intelligence problems. New work is focusing on adapting this technology to multisource data fusion for the Army.

5 Summary

Evidential reasoning has already been successfully applied to problems in several domains. However, the addition of the compatability relation to the theory of beliefs, the formalization and development of new evidential operators, and the use of graphical representations have greatly improved the overall usefulness and accessibility of these techniques.

References

[Coh85] Paul Cohen. *Heuristic Reasoning about Uncertainty: An Artificial Intelligence Approach.* Pitman Publishing, Inc., 1985.

[Dem68] Arthur P. Dempster. A generalization of Bayesian inference. *Journal of the Royal Statistical Society*, 30:205–247, 1968.

[DK77] R. Davis and J. J. King. An overview of production systems. In E. Elcock and D. Michie, editors, *Machine Intelligence 8*, pages 300–332, Ellis Horwood, Chichester, England, 1977.

[Doy81] Jon Doyle. A truth maintenance system. In Bonnie Lynn Webber and Nils J. Nilsson, editors, *Readings in Artificial Intelligence*, pages 496–516, Tioga Publishing Company, Palo Alto, California, 1981.

[GLF81] Thomas D. Garvey, John D. Lowrance, and Martin A. Fischler. An inference technique for integrating knowledge from disparate sources. In *Proceeding of the Seventh Joint Conference on Artificial Intelligence*, pages 319–325, American Association for Artificial Intelligence, 445 Burgess Drive, Menlo Park, California 94025, August 1981.

[Kon79] Kurt Konolige. Bayesian methods for updating probabilities. In R. O. Duda, P. E. Hart, K. Konolige, and R. Reboh, editors, *A Computer-Based Consultant for Mineral Exploration, Final Report, SRI Project 6415*, 333 Ravenswood Avenue, Menlo Park, California 94025, 1979.

[LB82] John F. Lemmer and Stephen W. Barth. Efficient minimum information updating for bayesian inferencing in expert systems. In *Proceedings of the National Conference on Artificial Intelligence*, pages 424–427, American Association for Artificial Intelligence, 445 Burgess Drive, Menlo Park, California 94025, August 1982.

[LG82] John D. Lowrance and Thomas D. Garvey. Evidential reasoning: a developing concept. In *Proceeding of the Internation Conference on Cybernetics and Society*, pages 6–9, Institute of Electrical and Electronical Engineers, October 1982.

[LG83a] John D. Lowrance and Thomas D. Garvey. *Evidential Reasoning: An Approach to the Simulation of a Weapons Operation Center.* Technical Report, Artificial Intelligence Center, SRI International, 333 Ravenswood Avenue, Menlo Park, California 94025, September 1983.

[LG83b] John D. Lowrance and Thomas D. Garvey. *Evidential Reasoning: An Implementation for Multisensor Integration.* Technical Report 307, Artificial Intelligence Center, SRI International, 333 Ravenswood Avenue, Menlo Park, California 94025, December 1983.

[Low78] John D. Lowrance. *Grasper 1.0 Reference Manual.* COINS Technical Report 78-20, Department of Computer and Information Science, University of Massachusetts, Amherst, Massachusetts 01003, December 1978.

[Low82] John D. Lowrance. *Dependency-Graph Models of Evidential Support.* PhD thesis, Department of Computer and Information Science, University of Massachusetts, Amherst, Massachusetts, September 1982.

[Low86] John D. Lowrance. *Grasper II Reference Manual.* Artificial Intelligence Center, SRI International, 333 Ravenswood Avenue, Menlo Park, California 94025, 1986. In preparation.

[LSG86] John D. Lowrance, Thomas M. Strat, and Thomas D. Garvey. *Application of Artificial Intelligence Techniques to Naval Intelligence Analysis.* Final Report SRI Contract 6486, Artificial Intelligence Center, SRI International, 333 Ravenswood Avenue, Menlo Park, California 94025, May 1986. In preparation.

[Sha76] Glenn Shafer. *A Mathematical Theory of Evidence.* Princeton University Press, Princeton, New Jersey, 1976.

[Sha86] Glenn Shafer. Belief functions and possibility measures. *The Analysis of Fuzzy Information*, 1, 1986. To appear.

[WA84] W. W. Wadge and E. A. Ashcroft. *Lucid, The Dataflow Programming Language.* Academic Press U. K., 1984.

[Wes86] Leonard P. Wesley. *Evidential Based Control in Knowledge-Based Systems.* PhD thesis, Department of Computer and Information Science, University of Massachusetts, Amherst, Massachusetts, 1986. In preparation.

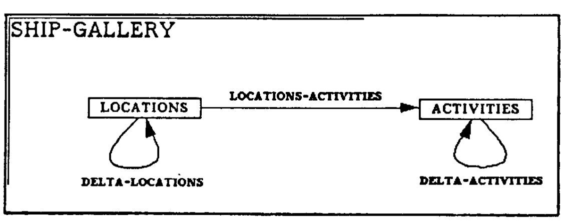

Figure 1: SHIP-GALLERY Gallery.

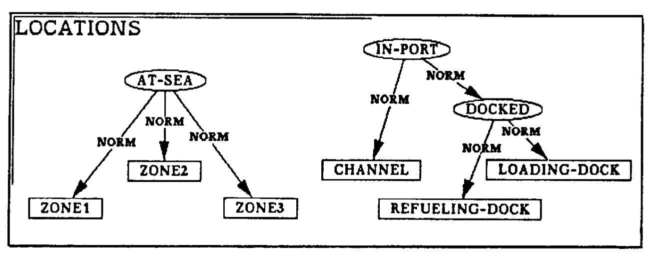

Figure 2: LOCATIONS Frame.

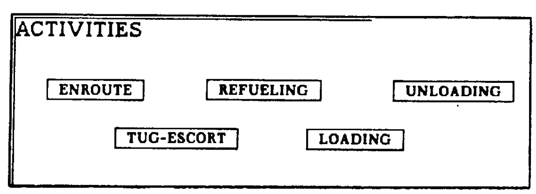

Figure 3: ACTIVITIES Frame.

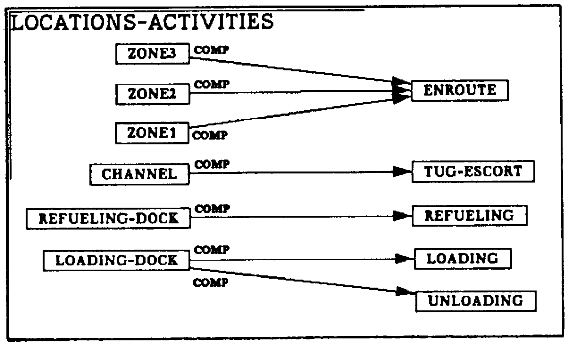

Figure 4: LOCATIONS-ACTIVITIES Compatibility Relation.

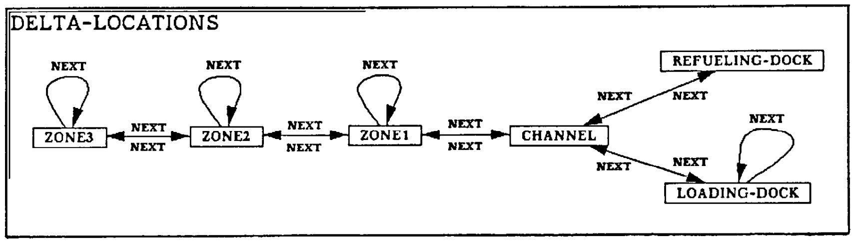

Figure 5: DELTA-LOCATIONS Compatibility Relation.

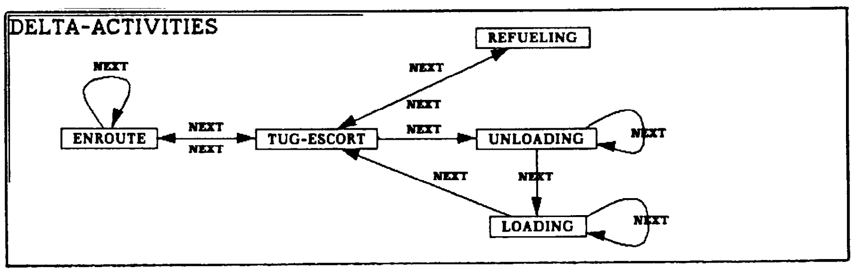

Figure 6: DELTA-ACTIVITIES Compatibility Relation.

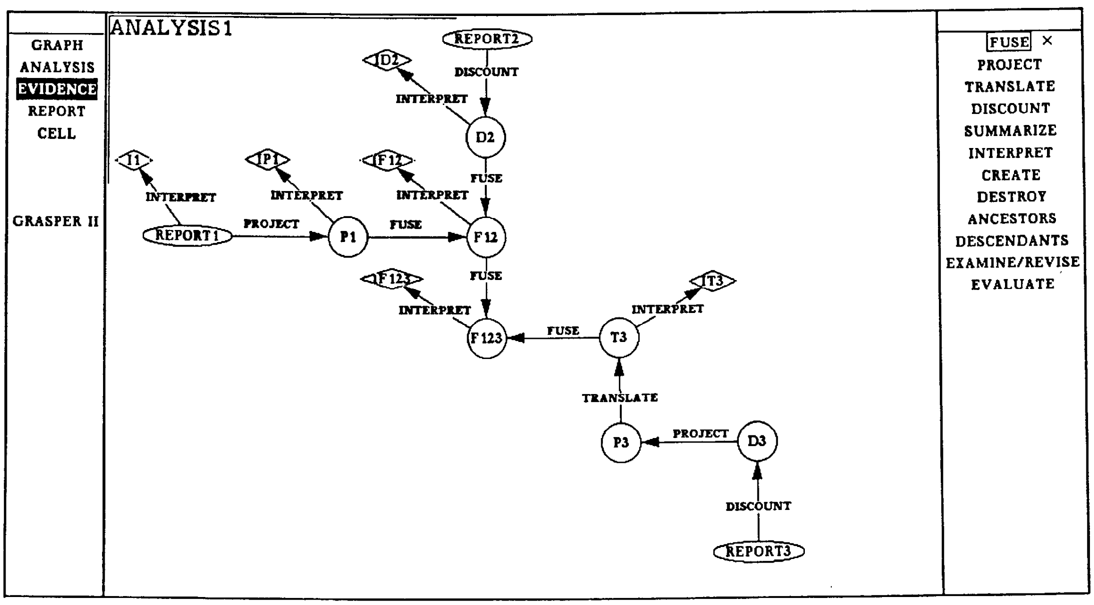

Figure 7: ANALYSIS1 Analysis.

```
REPORT1
TYPE: EVIDENCE
FOD: (LOCATIONS 1.)
MASSFUN: (((CHANNEL) 0.7) ((ZONE1) 0.3))
Exit □
```

```
P1
TYPE: PROJECTION
DELTA-T: 1.
FOD: (LOCATIONS 2.)
MASSFUN: (((REFUELING-DOCK LOADING-DOCK ZONE1) 0.7) ((ZONE2 CHANNEL ZONE1) 0.3))
Exit □
```

```
REPORT2
TYPE: EVIDENCE
FOD: (LOCATIONS 2.)
MASSFUN: (((CHANNEL LOADING-DOCK REFUELING-DOCK) 1.0))
Exit □
```

```
D2
TYPE: DISCOUNT
DISCOUNT-RATE: 20.
FOD: (LOCATIONS 2.)
MASSFUN: (((CHANNEL LOADING-DOCK REFUELING-DOCK) 0.8) ((REFUELING-DOCK ZONE2 CHANNEL LOADING-DOCK ZONE1 ZONE3) 0.2))
Exit □
```

```
F12
TYPE: FUSION
FOD: (LOCATIONS 2.)
MASSFUN: ((LOADING-DOCK REFUELING-DOCK) 0.56)
         ((CHANNEL) 0.24000001)
         ((REFUELING-DOCK LOADING-DOCK ZONE1) 0.14)
         ((ZONE2 CHANNEL ZONE1) 0.060000002))
CONFLICT: 0.0
Exit □
```

```
REPORT3
TYPE: EVIDENCE
FOD: (ACTIVITIES 3.)
MASSFUN: (((LOADING) 1.0))
Exit □
```

```
D3
TYPE: DISCOUNT
DISCOUNT-RATE: 40.
FOD: (ACTIVITIES 3.)
MASSFUN: (((LOADING) 0.6) ((TUG-ESCORT UNLOADING ENROUTE LOADING REFUELING) 0.4))
Exit □
```

```
P3
TYPE: PROJECTION
DELTA-T: -1.
FOD: (ACTIVITIES 2.)
MASSFUN: (((UNLOADING LOADING) 0.6) ((REFUELING LOADING ENROUTE UNLOADING TUG-ESCORT) 0.4))
Exit □
```

```
T3
TYPE: TRANSLATION
THETA: LOCATIONS
FOD: (LOCATIONS 2.)
MASSFUN: (((LOADING-DOCK) 0.6) ((ZONE3 ZONE2 ZONE1 LOADING-DOCK REFUELING-DOCK CHANNEL) 0.4))
Exit □
```

```
F123
TYPE: FUSION
FOD: (LOCATIONS 2.)
MASSFUN:((LOADING-DOCK) 0.5121951)
        ((LOADING-DOCK REFUELING-DOCK) 0.2731707)
        ((CHANNEL) 0.11707317)
        ((ZONE1 LOADING-DOCK REFUELING-DOCK) 0.06829268)
        ((ZONE2 ZONE1 CHANNEL) 0.029268293))
CONFLICT: 0.17999995
Exit □
```

Figure 8: Data from ANALYSIS1.