Computational Antiterrorism Solutions Using a General-Purpose Evidence Fusion Engine

Patrick J. Talbot and Dennis R. Ellis
TRW Systems, Command, Control and Intelligence Division

This paper describes our experience with the implementation of a general-purpose evidence fusion engine for assimilating data and predicting intent in the mission area of antiterrorism. Data fusion is a difficult, multifaceted problem with many unsolved challenges, the major one being that the data fusion problem is ill posed—data fusion means many things to many people. Fusion of evidence provides a focused solution in cases where a decision maker is faced with discrete bits of information that are inherently uncertain. The Dempster-Shafer Belief Network algorithm is the basis for our approach to data fusion. Advantages are that it provides intuitive results; differentiates belief, ignorance, and disbelief; and resolves conflicts. Insights into the behavior of the belief network were obtained from an analytical study. Mathematical and empirical properties were explored and codified. Belief networks have been constructed for seven mission areas, leading to the characterization of the prototype testbed as a general-purpose data fusion engine. A novel feature of the evolving testbed is that the user is allowed to override the belief or disbelief associated with a hypothesis and the network will self-adjust, or learn, the appropriate link values by instantiating the override. The back-propagation algorithm from artificial neural network research is used to adjust the links. Although the work described here is in the mission area of antiterrorism, it has broad applicability to decision making in any circumstance where evidence is uncertain, incomplete, possibly conflicting, and arrives asynchronously over time.

Introduction

This paper describes our experience with the tailoring of a general-purpose engine to the task of fusing evidence on terrorism, predicting enemy intent, and quantitatively comparing allied versus enemy operations. The tool provides semiautomated reasoning support for decision makers faced with risky decisions and unknown intents based on assessment information that is inherently uncertain, incomplete, and possibly conflicting. Decision makers often find it difficult to mentally combine evidence—the human tendency is to postpone risky decisions when data are incomplete, jump to conclusions, or refuse to consider conflicting data. Those versed in classical (Frequentist) statistics realize it does not apply in situations where evidence is sparse. A data fusion engine is needed.

Data fusion is a difficult, multifaceted problem with many unsolved challenges. A contributing factor to these challenges is that the problem is ill posed—data fusion means many things to many people. The taxonomy shown in Figure 1 attempts to organize and differentiate mathematical building blocks, correlation schemes, and true data fusion methods. Clearly, there are many types of data fusion, and the difference between fusion and correlation is tenuous. At the top are the fundamental building blocks, which are then

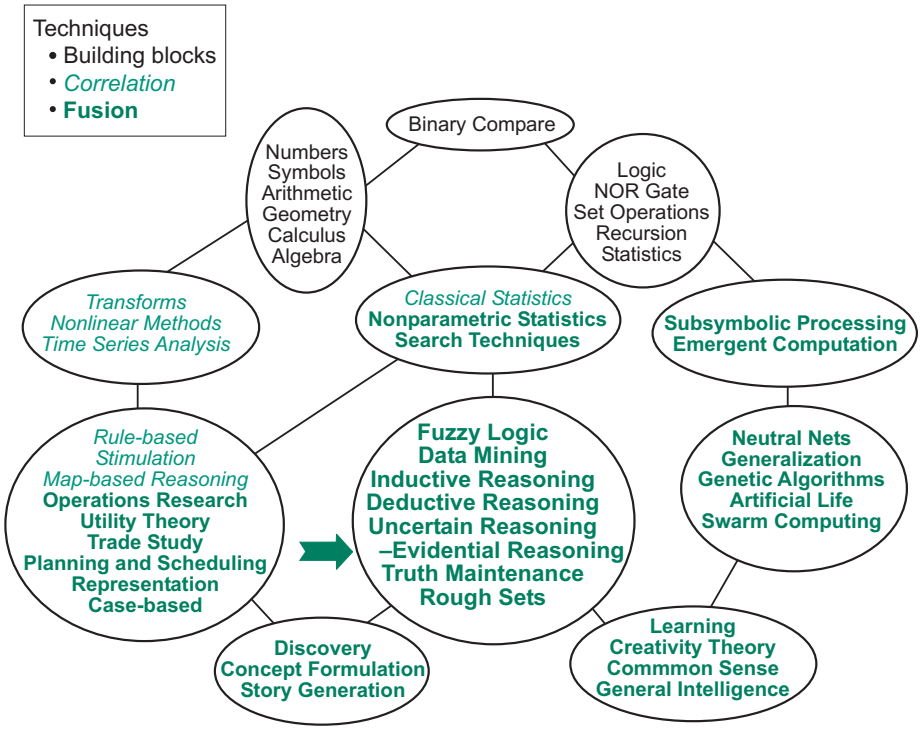

Figure 1. Fusion hierarchy

differentiated into analytical, statistical, and subsymbolic techniques. Uncertain reasoning, the topic of this paper, is among the statistical approaches.

Related techniques, such as Bayesian Networks and Rough Set Theory, were assessed for applicability. The evidential reasoning approach [1], which relies on the Dempster-Shafer Combination Rule, was chosen because it provides intuitive results, differentiates ignorance and disbelief (sometimes described as *skeptical* processing), and performs conflict resolution. Bayesian networks also provide intuitive results, but are better suited to causal reasoning. Rough sets differentiate between what is certain and what is possible and are of potential future interest—a truth maintenance system appears necessary to track the validity of hypotheses as evidence is accumulated.

Insights into the behavior of the belief network were obtained from an analytical study. Mathematical and empirical properties were explored and codified. Discovery of a representation for an identity and its inverse revealed fascinating properties with practical application; e.g., fusion equations are soluble in the same way as matrix equations.

A novel feature of the evolving testbed is that the user is allowed to override the belief or disbelief associated with a hypothesis and the belief network self-adjusts, or learns, the appropriate link values by instantiating the override. The back-propagation algorithm from artificial neural network research is used to adjust the links. In addition, the belief network constructs explanations of how outcomes are obtained. This is important in risky decision-making environments. The work described in this paper has broad applicability to risky decision making in circumstances where evidence is uncertain, incomplete, possibly conflicting, and arrives asynchronously over time.

Belief networks were constructed for seven mission areas, leading to the characterization of our prototype testbed as a general-purpose data fusion engine. The most sophisticated of these consisted of a pair of six-layer networks that mapped evidence and assessments to high-level terrorist and antiterrorist objectives to predict terrorist intent and derive best antiterrorism moves.

Approach

The Dempster-Shafer Combination Rule (Figure 2) for fusion of evidence is the core algorithm. Nodes in the network are represented as evidential intervals with values from the set of real numbers $[0 \leq n \leq 1]$. Three parameters specify each node: a *belief* (B), an *unknown* (U), and a *disbelief* (D). The words *unknown*, *ignorance*, and *don't know* are used interchangeably throughout this paper. The unknown parameter is computed as $U = 1–B–D$. New evidence is fused with existing evidence according to the equations in the figure. Although the theory of evidential reasoning, which is a skeptical brand of reasoning, allows n parameters within an evidential interval (producing computational complexity $2^n–1$), we found that specifying nodes with a single meaning and a set consisting of three values $\{B,U,D\}$ was satisfactory. *Disbelief* (D) is the complement of the more obtuse *plausibility* (P) parameter ($D = 1–P$) encountered in the literature.

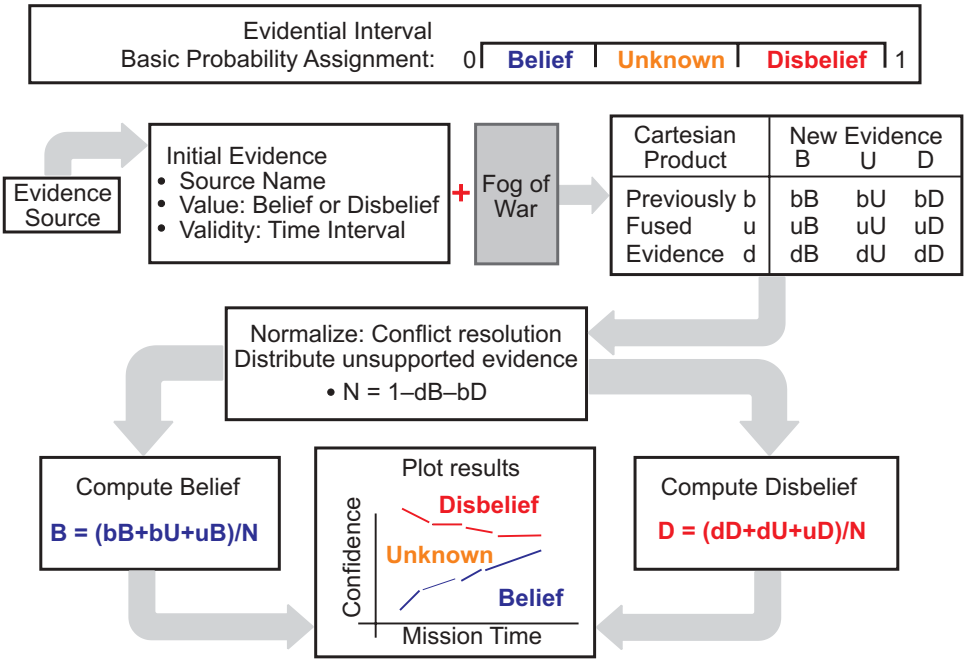

Figure 2. Dempster-Shafer Combination Rule

Having discussed how evidence is fused at nodes, we now discuss how nodes are combined with links to form a network. We use hierarchical directed a-cyclic graphs exclusively, since more complex representations are not required. The networks are composed of layers of nodes, each having values $\{B,U,D\}$, connected by links with constant values $\{L\}$. Three types of nodes form the networks:

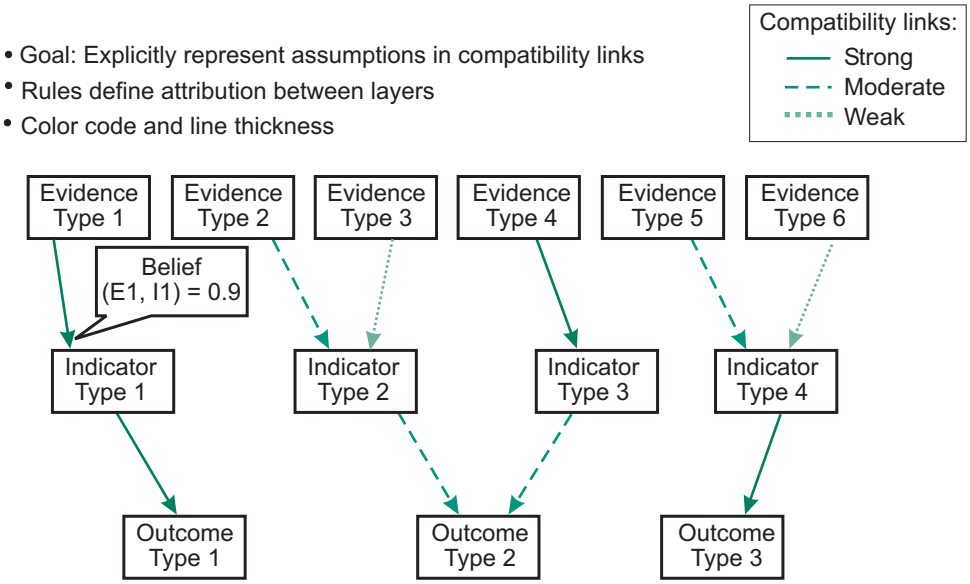

Figure 9. Belief network evolution

In addition, given two linked networks with the same topology, the effectiveness of one network as compared with the other can be shown. The result of the comparison is color coded based on the value of the ratio of the beliefs at the corresponding nodes. This is discussed in detail later.

Network Instantiations

Terrorist/antiterrorist. Terrorist attacks have historically come without warning. In retrospect, there have usually been enough shreds of evidence to suggest that the attack was about to occur, but these indicators reside in people's heads, on paper, and in multiple disparate databases. There has been no way to combine these various bits of information to provide timely warning. The initiative addressed here directly addresses the need for an integrated toolset to conceptually cluster information, fuse evidence, predict terrorist intent, and prioritize antiterrorist actions.

Evidence of terrorist activity is what drives the fusion engine. The content of this evidence is critically important to the formulation of the belief network—the representation on which the fusion engine works. We reasoned that a format for collecting evidence that captured answers to the reporter's questions—who, what, when, where, why, and how—as well as the original text was a good start (Figure 10). We took a "learn-by-doing" approach. A newspaper article [2] that contained about 100 pieces of evidence was the basis. We parsed each piece of evidence into the Reporter's Template, and concurrently added fields to accommodate practical aspects of the data set. Specifically, to produce a workable data structure, we added fields for *how severe* and *how certain*.

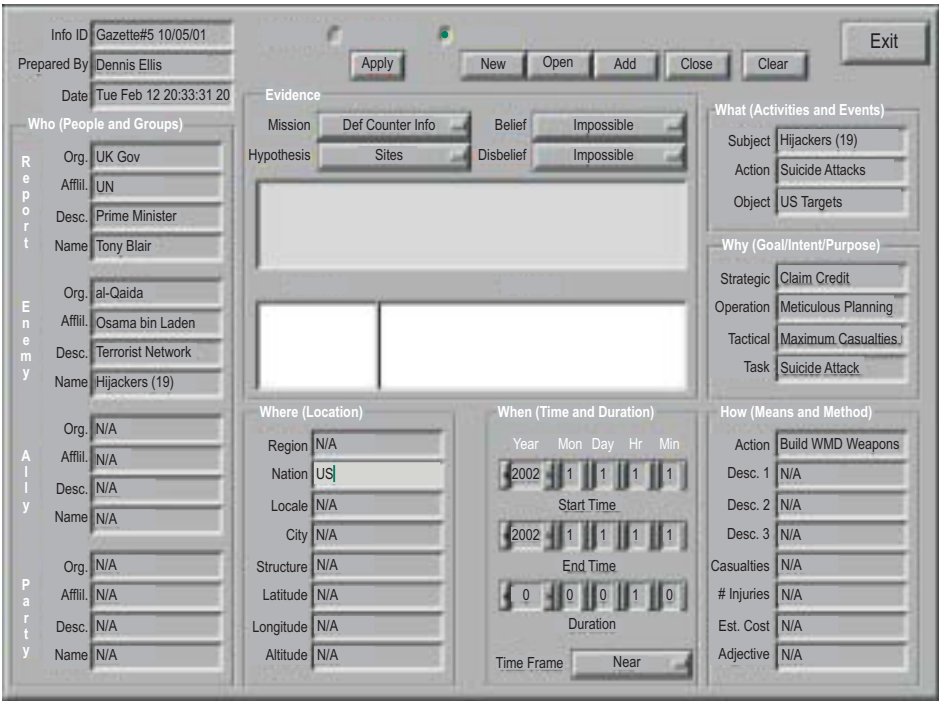

Figure 10. Reporter's questions template

Beyond providing a simple way to organize evidence of terrorist attacks and antiterrorism actions, the completed template proved useful in deriving a belief network and automatically producing the evidence to drive it. For example, the *Who* block indicates to which belief network the evidence relates. If *Threat* content is filled in, it goes on the terrorist belief network, while information in the *Friendly* section means it is associated with the antiterrorism belief network. Another finding was that the content of the *What* block provides events, activities, and tactical objectives for the belief network, while the *Why* block provides content for the operational objectives and strategic objectives in the belief network. The *When* block provides a key for chronologically sorting the evidence. Finally, the persistence entry is used in the belief network to degrade the belief in the data over time, if appropriate; e.g., weather predictions are not persistent but antiterrorist intents are.

The evidence templates were analyzed and became the basis for deriving the Terrorist Assessment belief network (Figure 11). Nodes are hypotheses, with the information sources row showing low-level hypotheses, the middle layers providing intermediate-level hypotheses, and the last row showing high-level objectives—also called intents. Although the algorithmic flow of evidence is from low-level input to high-level intent, most interest in this problem is predicting (low-level) terrorist events. We have also found that evidence may be introduced at any level in the hierarchy. To accommodate this reality, we extended the processing algorithm to allow evidence to be associated with any hypothesis. Evidence is then fused with existing evidence for the hypothesis.

The evidence templates and the terrorist belief network formed the basis for the antiterrorism belief network (Figure 12). In many cases, our activities, tactical objectives, and

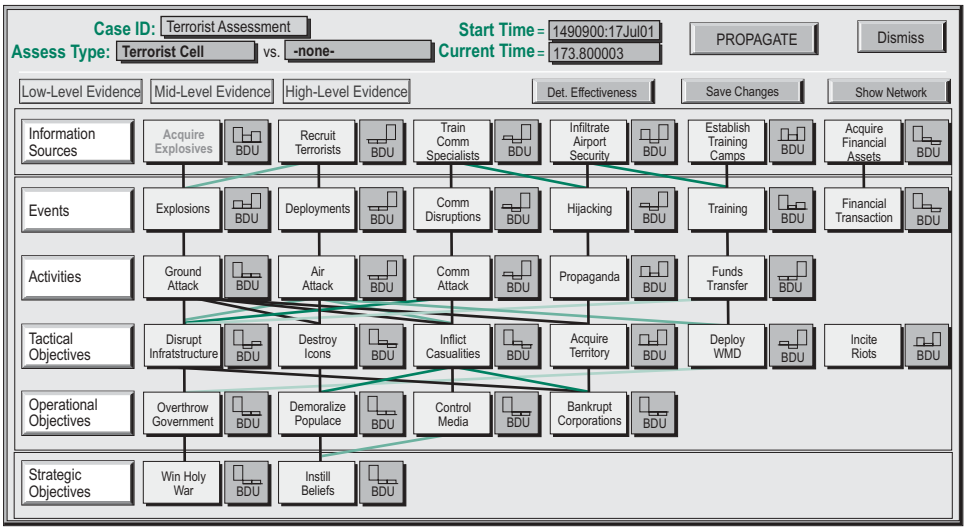

Figure 11. Terrorist assessment belief network

strategic objectives related directly to antiterrorist nodes. In other cases, asymmetries were evident. For example, air superiority is an antiterrorist operational objective that is virtually unchallenged by terrorists.

The antiterrorism and terrorist belief networks were then compared. The question was "What strengths and weaknesses are the terrorists or the antiterrorism forces likely to exploit?" Algorithmically, we compute the ratio of antiterrorism belief to terrorist belief for each node and use it as a score. A ratio of unity is defined as parity, while a ratio greater than 10 corresponds to a significant antiterrorism strength, and a ratio less than 1.0 corresponds to a terrorist strength. A ratio is tabulated for each node, and portrayed using color coding on an effectiveness display (Figure 13). This effectiveness algorithm is analogous to assessing the relative worth of chess moves using a (one ply) minimax algorithm, although our implementation is much simpler.

Other Examples. In addition to the terrorist/antiterrorist instantiation, we have tested the belief network in a number of other mission areas to show its versatility. These include

- Combat assessment for effects-based operations
- Weather prediction effects
- Kill chain belief network
- Space weather fusion
- Foreign missile launch assessment (with and without fog-of-war effects)
- Hazardous material alert

Combat Assessment. How do we fuse battle damage assessment reports, intelligence information, and operational knowledge over a period of weeks to determine if strategic, operational, and tactical objectives are being met? The combat assessment instantiation shows how weapon-target-assessment information influences belief about task success, which in turn influences beliefs about tactical, operational, and strategic objectives.

Weather Effect Predictions. How will the 24-hour forecast affect tomorrow's Air Tasking Order for sorties that employ laser-guided munitions? This instantiation shows if weather conditions obscure an optical path between a laser designator and a target.

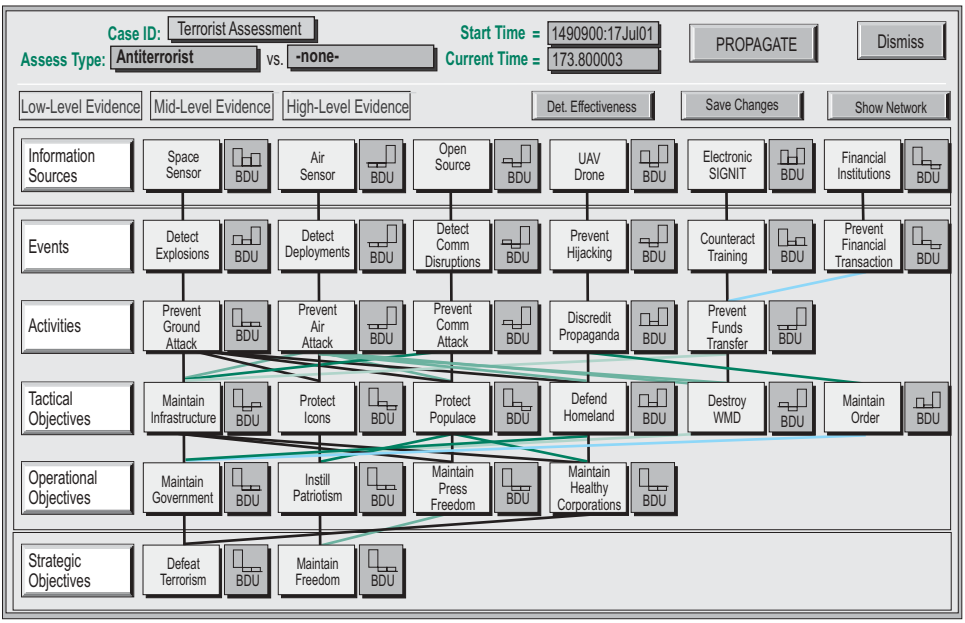

Figure 12. Antiterrorism assessment belief network

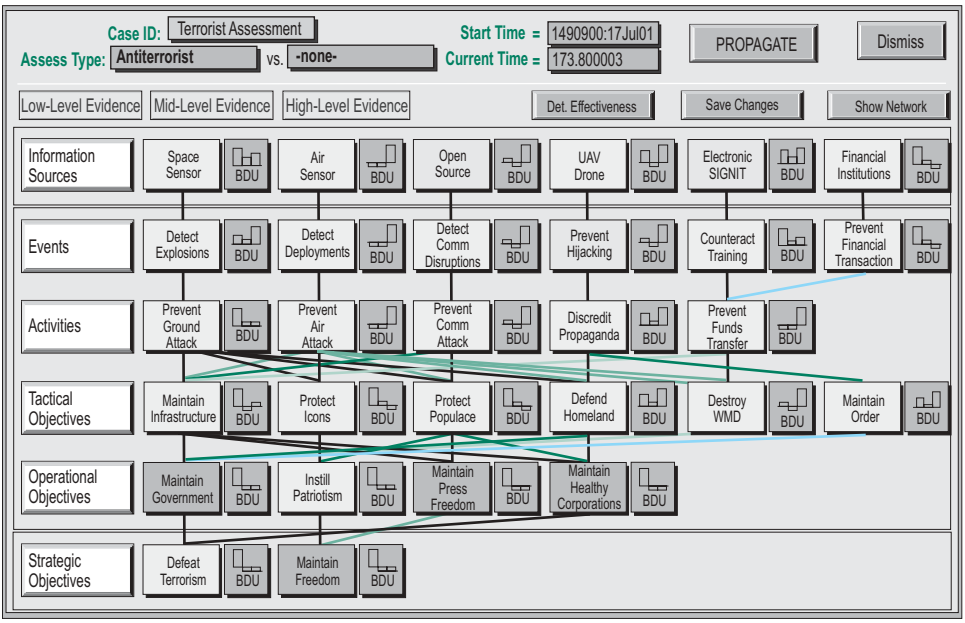

Figure 13. Effectiveness of antiterrorism assessment versus terrorist assessment

Kill Chain Belief Network. For time critical targets (TCTs), such as mobile missile launchers, a variety of information sources provide information on how well targets are found⇒fixed⇒tracked⇒targeted⇒executed⇒assessed (referred to as the kill chain). This instantiation allows TCT kill-chain evidence to be posted, organized, and coherently represented in real time.

Space Weather Fusion. Over a period of hours to days, a satellite control officer may be confronted by a bewildering variety of information, including communications outages, reports of ionospheric scintillation, radio frequency interference reports, and indicators of threat activity. This instantiation allows incomplete evidence to be fused.

Foreign Launch Assessment. This early demonstrator was based on the need to fuse assessments from space sensors, radar, and intelligence reports to determine, as a function of time, whether a possible foreign missile launch was hostile, deliberate, and whether it was an intercontinental ballistic missile (ICBM) or a spacecraft launch.

Hazardous Material Alert. We produced and demonstrated a belief network to Northern Command. In the underlying scenario, evidence that arrived from five sources with various beliefs was related to intermediate indicators such as crowded hospitals and fallout. We fused this evidence to determine whether an accident, hostile activity, or natural phenomena occurred.

Fog-of-War Effects. A software module that perturbs data and human cognitive performance was integrated with [3] the Foreign Launch Assessment belief network. Examples of the factors that we allow the user to specify are the frequency and severity of lost data, data overload, human confusion, and false assumptions, among others. We found that the result of applying Fog-of-War perturbations to the evidence is significant.

Current Work

We are investigating the use of belief networks to tell a story (as represented by a belief network) and to operate on these networks as though they were mathematical functions. Stories are constructed by associating characteristics with each hypothesis (node) in the network. The characteristics provide answers to the *who*, *what*, *where*, *when*, *why*, *how*, *how much*, and *how certain*. Mathematical operations on networks that we are investigating include differentiation to provide a sensitivity analysis showing the impact of a perturbation of a node on other nodes, integration to determine best paths through a network, a metric that defines the level of belief in the story as a whole, intersections and unions of networks to expand on stories, and recursion to allow nesting of networks. A TRW-sponsored workshop, *Computer Algorithms for Human Decision Making*, was conducted May 2002. Sessions included "Uncertain Reasoning," "Multistrategy Reasoning and Planning," "Information Retrieval," and "Innovative Applications." Proceedings are available from the authors.

Summary

The work described in this paper has broad applicability to decision making in circumstances where evidence is uncertain, incomplete, possibly conflicting, and arrives asynchronously over time. The Dempster-Shafer theory of evidence has been found to have very useful mathematical properties. In particular, an inverse has been discovered that allows fusion equations to be solved arithmetically. In addition, the derived properties of these belief networks collectively suggest intuitive application of the technique as a general-purpose fusion engine for uncertain reasoning. A novel feature of our implementation is the addition of a back-propagation algorithm that allows the user to override beliefs

and disbeliefs in nodes. The back-propagation algorithm adjusts precursor node and link values to reconcile the network. Thus, the network learns from user training data in the form of overrides.

References

1. D. Zarley et al., "Evidential Reasoning Using DELIEF," *Readings in Uncertain Reasoning*, Morgan Kauffman, San Mateo, California, 1990.
2. "The Case Against Osama bin Laden," Colorado Springs Gazette, October 5, 2001, which contained the full text of Tony Blair, Prime Minister of the United Kingdom.
3. P. Talbot, "Simulating Fog-of-War in Military Decision Making," presented at International Conference on Industrial and Engineering Applications of Artificial Intelligence and Expert Systems (IEA/AIE 2000), New Orleans, Louisiana, July 2000. The paper was not included in the Proceedings.

Patrick J. Talbot has extensive experience with TRW, most of which is in Independent Research and Development (IR&D) efforts. He is currently the principal investigator on the Information Analysis and Planning IR&D project. He has performed on Minuteman, spacecraft projects, battle management tasks, antisatellite systems, and the development of decision algorithms. In addition, he has managed a Delivery Order titled "Technology Insertion Studies and Analysis" at the Joint National Test Facility in Colorado Springs, Colorado, where his research included parallel computing, simulated human behavior, and strategic planning. He holds a BS in applied mathematics, a BS and MS in physics, all from Pennsylvania State University.

pat.talbot@trw.com

Dennis R. Ellis was the deputy lead of the Technology Insertion Studies and Analysis project at the JNTF. He worked in the fields of parallel computing, performance benchmarking, simulated human behavior, and strategic planning. Previously, he was assistant manager in the Marketing and Benchmark group at Cray Research. Since joining TRW five years ago, he has worked in missile defense and the Information Analysis and Planning IR&D project. He holds a BS in chemical engineering, from the Colorado School of Mines.

dennis.ellis@trw.com