DATA FUSION

Data fusion consists of a set of quantitative and qualitative modeling techniques for integrating the reports of multiple, diverse sensors for the purpose of modeling targets in a domain of interest. These techniques must address the aggregation of selected sensor reports into "tracks" of hypothesized targets, estimate the current and predict the future positions of these tracks, infer the identification of the track in terms relevant to the domain, and begin higher level inferences related to the current status of the situation comprising the entire set of tracks and then consider future possibilities given the preceding predictions and inferences.

OVERVIEW

The basis of all fusion processes is the synergistic use of redundancy and diversity of information contained in multiple, overlapping observations of a domain to achieve a combined view that is better than any of the individual observations. The fusion process may improve a number of performance metrics over that which is achievable by any single source or sensor, for example, accuracy, resolution, timeliness, or state estimates (of individual entities or events). The fusion process may also expand the understanding of relationships between entities or events and the overall comprehension (of the entire domain). Finally, the fusion process may expand the spatial domain covered by that available to any single sensor (1–5).

The data fusion process has also been referred to as multisensor fusion or sensor fusion (used for real-time sensor system applications), multisource fusion (referring to intelligence and law enforcement applications that combine intelligence sources), sensor blending, or information fusion. Data fusion generally refers to automated processes. However, manual data fusion processes have long been performed by humans to process volumes of data in numerous applications, for example, detective work in law enforcement, weather analysis and prediction, statistical estimation, intelligence analysis, and air traffic control, to name a few. Animals perform neurological data fusion processes by combining sensory stimuli and applying cognitive processes to perceive the environment about them. By combining sight, sound, smell, and touch, humans routinely reason about their local environment, identifying objects, detecting dangerous situations, while planning a route.

The U.S. DoD Joint Directors of Laboratories established a data fusion subpanel in the mid-1980s to establish a reference process model and a common set of terms for the functions of data fusion. This model (Fig. 1) defines four levels, or stages, of functions that are oriented toward intelligence and military applications.

The first level is object refinement, where sensor or source reports that contain observations (detections of objects and measurements about those objects) are aligned, associated, and combined to refine the estimate of state (location and kinematic derivatives) of detected objects using all available data on each object. The sequence of operations within object refinement are as follows.

1. *Alignment.* All observations must be aligned to a common spatial reference frame and a common time frame. For moving objects, observed at different times, the trajectory must be estimated (tracked) and observations propagated forward (or backward) in time to common observation time.
2. *Association.* Once in a common time-space reference, a correlation function is applied to observations to determine which observations have their source in the same objects. Correlation metrics generally include spatial (same location), spectral (similar observed characteristics) and temporal (same time of appearance) parameters. If the correlation metric for a pair of observations is sufficiently high, the observations are assigned to a common source object.
3. *State Estimation.* The state of the object (the location if the object is stationary, or a dynamic track if the object is moving) is updated using all associated observations.
4. *Object Identification.* The identity of the object is also estimated using all available measurements. If the sensors measure diverse characteristics of the object (e.g.,

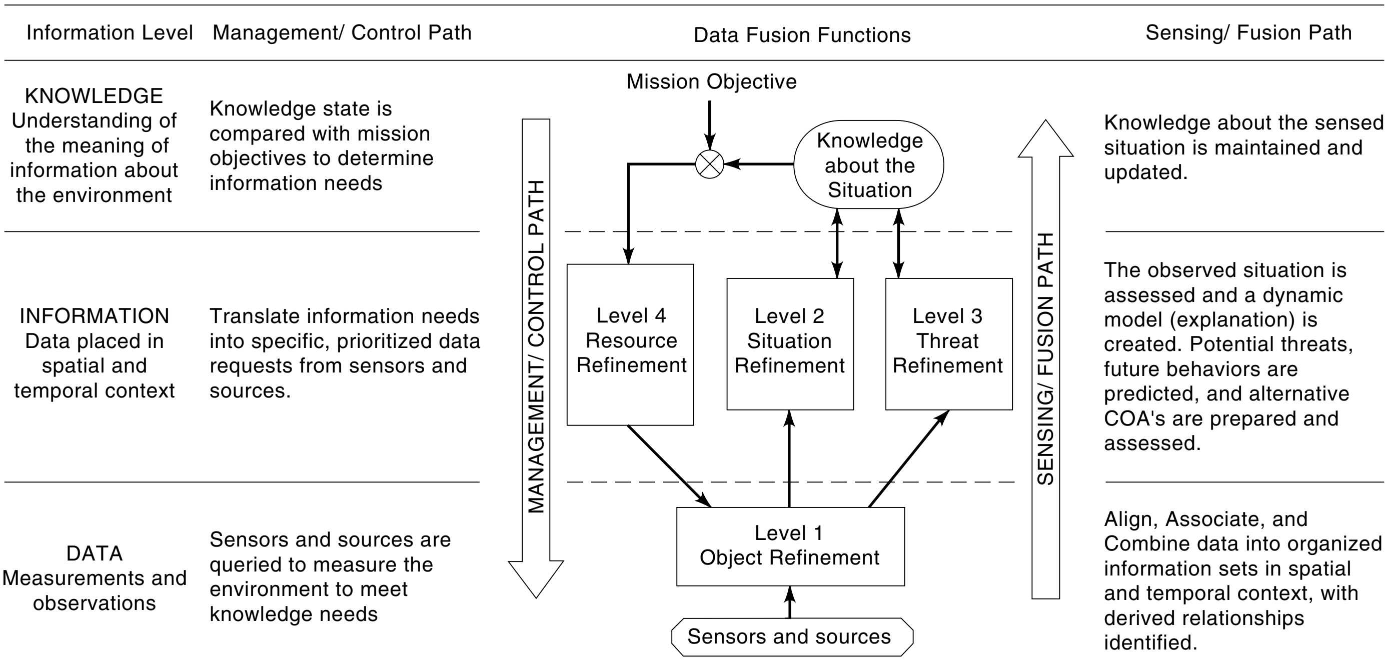

Figure 1. Four levels of functions oriented toward intelligence and military applications.

color and shape) automatic classification techniques are applied to resolve the identity.

The inputs to object refinement are reports containing raw data, and the output product is organized and refined data, or information.

The next levels of processing attempt to understand the information—to create knowledge about the observed objects and their behavior as groups in their context. The second level is situation refinement, which has the goal of understanding the meaning of the assembly of objects by detecting relationships between objects, detecting aggregate sets of objects, and identifying patterns of behavior to create a model of the current situation—a scene. The third level is threat refinement, which looks to the future to predict potential courses of action (COAs) of objects and groups within the situation scene that may pose a threat (defensive focus) or opportunity (offensive focus) for action by the user of the fusion system.

The fourth level is process refinement, which governs the overall fusion process to optimize the use of data to achieve the knowledge objectives of the system. Sensors are managed, and internal processes are adapted to optimize the information and knowledge products produced by the process.

MAJOR ISSUES AND ALGORITHMS

This section addresses the major topics of data fusion in more detail. For object refinement, the three most algorithmic intensive processes are association, state estimation, and object identification. Alignment is a measurement intensive process. Situation and threat refinement employ higher level reasoning techniques from Artificial Intelligence. The section below on situation and threat refinement addresses many of the techniques relevant to these two phases. Resource refinement is relatively new, and a cohesive body of literature is still developing, so this area is not discussed below. Techniques commonly discussed for resource refinement are knowledge-based systems, maximizing entropy, and decision analysis.

Association of Sensor Reports

The purpose of data association is to aggregate individual sensor reports from a common target into a group and to develop a dynamic model that represents a "track" based on the sequence of reports in the group. Generally, an initial set of sensor reports are collected in a "batch," and "report to report" association is used to establish the first set of tracks. Report to report association must continue to examine the possibility of finding new targets and to create the associated tracks. Once tracks are established, it is common to perform "report to track" association, that is, to determine if new sensor reports can be associated with existing tracks to assume that these reports are being received from sensed emanations of the hypothesized target. Multisensor architectures can either have one centralized track to which all sensor reports are associated or can form tracks for each sensor and then perform "track to track" association. The centralized track structure is generally considered the optimal approach. The association decision process is based on the notion of applying a correlation test to determine if pairs of reports, or a report and a track, are "close enough" to be associated.

The optimum solution to the data association problem is a Bayesian solution because the problem is characterized by the randomness of target motion and sensor observations. Unfortunately, the Bayesian solution has a number of implementation problems that inhibit its success. First, the Bayesian solution requires prior probabilities about the number of targets, their locations, and their identity. Information for the formation of these prior probabilities is often available but is substantially more subjective than the uncertainties concerning sensor capabilities and, to some degree, target motion. The second requirement for the Bayesian solution are the relative likelihoods of sensor reports given targets; this is the one input that is generally available and used in all solution approaches. Finally the Bayesian solution requires maintaining a multiple hypothesis inventory over many sensor reporting periods of an exponentially growing number of hypothetical associations and therefore hypothetical tracks.

Data association techniques can be characterized as deferred logic (or multiple hypothesis) techniques and sequential assignment techniques. The deferred logic techniques attempt to address the complete solution space addressed by the optimal Bayesian method but usually opt for a maximum likelihood criterion that is consistent with assuming the prior probability distributions over spatial location and number of the targets are uniform. Deferred logic techniques emphasize pruning relatively low probability hypotheses and merging hypotheses that are similar in space and identity in order to handle the exponential growth of hypotheses (4,6,7).

Sequential assignment techniques either formulate the problem as a hard assignment of each report to a report or to a track (or to a false alarm), which is the generalized assignment problem, or as a virtual report assignment, which involves the probabilistic computation of relative likelihoods (6–8). In the hard assignment approach (the most commonly used approach), the goal is to assign recently arrived sensor reports to existing tracks or other recent reports or discard them as false alarms such that the combined likelihood of these assignments is maximized. The likelihood of each assignment is often measured by an inverse function of the "closeness" between the report and the track or report under consideration for assignment; maximum likelihood is therefore also minimum distance. Solution approaches to the assignment problem consist of Lagrangian relaxation, relaxation for network flows, a generalization of the Signature method, and an auction algorithm for the transportation problem (8). A number of good, suboptimal assignment algorithms have been developed: the Munkres, Ford-Fulkerson, and the Hungarian methods (6). The most often implemented approach is called the "greedy" or "row-column heuristic" assignment algorithm in which a track is picked at random, the reports from each sensor that are closest to it are assigned to it, a second track is picked at random, the remaining reports from each sensor that are closest to it are assigned it, and so on. This assignment approach is particularly suited to the difficult problem of report to report association for passive sensors. Because passive sensors do not provide any range information, the problem of an exponentially growing number of ghost targets arises as the number of sensors and reports increases (8). This passive sensor association problem is NP ("nondeterministic polynomial") hard (i.e., in all likelihood cannot be solved by an algorithm of polynomial time complexity) when the number of sensors is three or greater; the deferred logic approach becomes completely impractical for this problem when there are 10 or more targets.

The other sequential assignment technique is called Joint Probabilistic Data Association (JPDA) (4,6,7). JPDA creates virtual reports by combining the likelihoods that all assignable reports belong to an existing track or should be used to create a new track. This approach has been used successful when the number of targets is known and is small compared to the clutter (false alarm) environment.

Table 1. Common Distance Measures Between Vectors

Distance Measure Name	Distance Measure
City Block	$\lvert x_1 - x_2 \rvert$
Euclidean	$[(x_1 - x_2)^2]^{1/2}$
Minkowski	$[(x_1 - x_2)^m]^{1/m}$
Weighted Euclidean	$[(x_1 - x_2)^T W (x_1 - x_2)]^{1/2}$
Mahalanobis	$[(x_1 - x_2)^T R^{-1} (x_1 - x_2)]$
Bhattacharyya	$2[(x_1 - x_2)^T (R_1 + R_2)^{-1} (x_1 - x_2)] + 4 \ln [\lvert (R_1 + R_2)/2 \rvert / (\lvert R_1 \rvert^{1/2} \lvert R_2 \rvert^{1/2})]$
Chernoff	$(\alpha(1 - \alpha)/2)[(x_1 - x_2)^T ((\alpha R_1 + (1 - \alpha) R_2)/2)^{-1} (x_1 - x_2)] + 0.5 \ln [\lvert (\alpha R_1 + (1 - \alpha) R_2)/2 \rvert / (\lvert R_1 \rvert^{\alpha} \lvert R_2 \rvert^{1-\alpha})]$
Divergence	$0.5[[(x_1 - x_2)^T (R_1 + R_2)^{-1} (x_1 - x_2)] + \mathrm{tr}(R_1^{-1} R_2 + R_2^{-1} R_1 - 2I)]$
Product	$2[(x_1 - x_2)^T (R_1 + R_2)^{-1} (x_1 - x_2)] + 2 \ln [4\pi^2 \lvert (R_1 + R_2) \rvert]$

There are many possible dimensions on which to measure the closeness of reports to each other or the closeness of reports to tracks: spatial, velocity, and acceleration dimensions; target characteristics such as size, shape, and area; target emanations such as radar, ultraviolet, and infrared. Unfortunately, reports from different sensors will not contain the same dimensions. For example, a radar may provide two spatial and velocity dimensions (range and azimuth), while an infrared sensor provides the spatial dimensions of azimuth and elevation angles. The data association must be performed on the overlap of dimensions; this overlap may be very limiting. The more sensors there are contributing to the data association problem, the greater the overall pool of overlapping dimensions.

Nine of the most common distance measures (correlation metrics) used in association are shown in Table 1. These measures define the distance between two vectors x_1 and x_2 that represent the overlap of dimensions available from a sensor report and an existing track (or other sensor report). The city block distance measures distance along the sides of the rectangle formed by the measurements. The Euclidean distance is the shortest distance between two points. The Minkowski distance is the generalized Euclidean distance, where m is a number greater than 1 and not equal to 2. The Euclidean distance can also be generalized by weighting the dimensions on some basis to reflect that fact that the dimensions might not be commensurate, for example, spatial distance and velocity. The Mahalanobis distance is a special weighted generalization in which the weighting matrix is the covariance matrix (R) for the vectors representing the uncertainty associated with the measurement process; both vectors are assumed to have equal covariance matrices which work for report to report association when the reports are from the same sensor. Bhattacharyya further generalized the Mahalanobis distance by allowing unequal covariance matrices for the two vectors, a covariance for the measurement process error of each sensor or for the sensor and the track. Chernoff's definition of distance is a further generalization of the Bhattacharyya distance. The Divergence and Product measures have also been proposed.

State Estimation

In data fusion, the kinematic state of the track is typically estimated by a tracking algorithm known as the Kalman filter. The filter estimates the kinematic state of the target object using a sequence of all reports associated with that target (6–7). The development of the Kalman filter assumes a motion model for the track and a measurement model for the sensors. The typical form of the motion model is

$$\boldsymbol{x}(k+1) = \Phi\boldsymbol{x}(k) + \boldsymbol{q}(k)$$

where $\boldsymbol{x}$ is the n-dimensional state vector of the target, spatial position, velocity, etc.;
k is the discrete time increment;
Φ is the transition matrix from one state to another;
$\boldsymbol{q}$ is the n-dimensional noise vector for target motion, called plant noise;
$\boldsymbol{Q}$ is the n by n covariance matrix of the plant noise.

The measurement model usually takes the form

$$\boldsymbol{y}(k) = \boldsymbol{H}\boldsymbol{x}(k) + \boldsymbol{v}(k)$$

where $\boldsymbol{y}$ is the m-dimensional measurement vector of $\boldsymbol{x}$;
$\boldsymbol{H}$ is the m by n measurement matrix;
$\boldsymbol{v}$ is the measurement noise;
$\boldsymbol{R}$ is the covariance matrix of the measurement noise.

Note, for multisensor systems, there will be a measurement model for each sensor. Also, note that the covariance matrices associated with plant and measurement noise may vary with time.

The Kalman filter equations become:

$$\begin{aligned}
\hat{\boldsymbol{x}}(k|k) &= \hat{\boldsymbol{x}}(k|k-1) + \boldsymbol{K}(k)[\boldsymbol{y}(k) - \boldsymbol{H}\hat{\boldsymbol{x}}(k|k-1)]\\
\boldsymbol{K}(k) &= \boldsymbol{P}(k|k-1)\boldsymbol{H}^T[\boldsymbol{H}\boldsymbol{P}(k|k-1)\boldsymbol{H}^T + \boldsymbol{R}]^{-1}\\
\boldsymbol{P}(k|k) &= [I - \boldsymbol{K}(k)\boldsymbol{H}]\boldsymbol{P}(k|k-1)\\
\hat{\boldsymbol{x}}(k+1|k) &= \Phi\hat{\boldsymbol{x}}(k|k)\\
\hat{\boldsymbol{P}}(k+1|k) &= \Phi\boldsymbol{P}(k|k)\Phi^T + \boldsymbol{Q}
\end{aligned}$$

where $\hat{\boldsymbol{x}}(k|k)$ is the estimate of $\boldsymbol{x}$ at time k given the sensor data through time k.

A major implementation issue associated with the Kalman filter is partitioning the state vector into independent subvectors so that multiple, simpler Kalman filters can be used to compute the estimated state at each point in time. Significant work has taken place on multiple sensor Kalman filtering and the simultaneous use of multiple Kalman filters based on different motion models to derive the best estimate of the target state in situations in which the target may turn, dive, or climb, and accelerate or decelerate.

The Kalman filter assumes there is one, linear motion model of the target's motion and a linear measurement model. Enhancements of the Kalman filter are called the extended Kalman filter (EKF) and the interacting multiple model (IMM) Kalman filter. The EKF is a suboptimal estimation algorithm that assumes nonlinear dynamic or measurement models. The IMM posits several motion models for the target. The IMM mixes hypotheses across the filter results from the previous time period in developing an overall estimate in the current time period, which is used to initialize the several filters for the next time period (7).

Object Identification

The identification of objects that are being tracked on the basis of multiple sensor reports must deal with the fact that the sensors are again observing the same object using different phenomonology, and possibly from different vantage points. Most targets are not completely symmetrical along either vertical or horizontal axes. The advantages of being able to "see" a target through different "lenses" and different perspectives also introduces a number of complications. These complications are ultimately reflected in uncertainty about what was seen and what it means to have seen that. As a result, most methods for solving the object identification problem are based upon modeling and updating uncertainty using new information contained in the sensor reports.

There are hard sensor integration techniques that rely on each sensor declaring what type of target the sensor believes the object to be. These techniques then use rule-based algorithms that address the combination of the different sensor declarations with the ability of the sensors to make accurate identification declarations. In complex environments with many types of objects, these techniques tend to fail often due to a brittleness in their logic structure and the limited resolution of their representation of uncertainty.

The three most commonly proposed methods for dealing with uncertainty explicitly are probability theory, evidence theory, and fuzzy sets.

Probability Theory. Probability theory is the common method for representing uncertainty for object identification (1,9–10). Within probability theory, Bayes's rule addresses how one should optimally update one's uncertainty as new information becomes available.

$$p(ID_k|rpts, \text{prior info}) = \frac{p(rpts|ID_k, \text{prior info})p(ID_k|\text{prior info})}{p(rpts|\text{prior info})}$$

Bayes's rule says the posterior probability of the kth ID being correct given a set of sensor reports and the prior information about the situation equals the likelihood of receiving those reports given the kth ID being correct and the prior information times the probability of the kth ID being correct given the prior information, divided by the probability of the reports being received given the prior information. A common criticism of the Bayesian approach is that the prior information is only available from expert judgments; yet these experts commonly have a great deal of information about the objects in the environment and the characteristics of these objects. The likelihood of specific reports given the various possible identifications typically requires much of the same information, and these likelihoods are crucial for the hard sensor integration techniques as well as many of the competing approaches.

Evidence Theory. Evidence theory operates on a frame of discernment, Θ, for which there are a finite number of focal elements. The power set of Θ, 2^Θ, is the set of all subsets of Θ. Evidence theory allows one to attach a probability to any member of the power set of the frame of discernment. Evidence theory (1) uses Dempster's rule to combine multiple belief functions, say from different sensors, and then (2) computes the supportability and plausibility measures for each element of the power set of Θ. Dempster's rule and the measures called supportability and plausibility are defined as follows (11–12).

A valid belief function assigns a probability measure to each element of 2^Θ. It is easily shown that for a frame of discernment with n focal elements, there will be 2^n elements in the power set, including the null set and Θ itself. Evidence theorists interpret a probability greater than 0 that has been assigned to Θ as the probability that could be assigned to any of the other $2^n - 1$ elements of the power set. This probability is said to be "uncommitted" and is described as a measure of ignorance. This is the concept that is truly unique to evidence theory and, as shown later, causes the two algorithms to diverge. One major effect of uncommitted belief is: $b(A) + b(A') \leq 1$, where $b(A)$ is the probability assigned to A which is an element of 2^Θ, and $b(A')$ is the probability assigned to all sets of 2^Θ whose intersection with A is null.

Since Θ has a non-null intersection with A and A', any time the belief in $\Theta > 0$, the above equation will be a strict inequality.

Dempster's rule is defined as follows for any two belief functions $b_1(A)$ and $b_2(B)$, where Θ has been divided into two possibly different representations $\{A_n\}$ and $\{B_m\}$:

$$b(A_i \cap B_j) = \frac{b_1(A_i)b_2(B_j)}{1-Q}$$

where

$$Q = \sum_r \sum_s b_1(A_r)b_2(B_s)$$

such that $A_r \cap B_s$ are all instances of the null event.

It is easily shown that Dempster's rule is equivalent to Bayes's rule when there is no uncommitted belief. It can also be shown that the uncommitted belief will decrease with the addition of each new belief function having some committed belief.

Since evidence theory allows some belief to be uncommitted, it is possible to develop lower and upper measures of uncertainty for any element of 2^Θ. Supportability, the lower measure, is defined by

$$s(A_j) = \sum_r b(A_r), \quad \forall A_r \subseteq A_j$$

The upper measure, called plausibility, is

$$pl(A_j) = 1 - s(A'_j), A'_j = \bigcup_s A_{sr} \quad \forall A_s \ni A_s \cap A_j = \varphi$$

Fuzzy Sets. Another approach to integrating various statements associated with ambiguous measures of uncertainty is fuzzy sets (13). In fuzzy sets, the membership function of a set is allowed to take on any value in the closed interval from 0 to 1:

$$z_\Psi(x) \colon \Omega \to [0, 1]$$

where z is the membership (or characteristic) function
Ψ is the fuzzy set
x is a focal element of the fuzzy set
Ω is the universal set.

Table 2. TWS Sensor Reports

Report States	TWS Report 1	TWS Report 2
SAM-X (any state)	0.0	0.2
SAM-Y (any state)	0.3	0.2
SAM-Y.ttr	0.0	0.4
SAM-Y.acq or SAM-Y.ttr	0.2	0.0
SAM-Y.ttr or SAM-Y.ml	0.4	0.0
Unknown (uncommitted)	0.1	0.2

From this definition and a number of axioms (e.g., commutative and associative), the most common fuzzy operators on fuzzy set operations are

$$z_{\overline{\Psi}}(x) = 1 - z_\Psi(x) \forall x \ni z_\Psi(x) > 0$$
$$z_{\Psi \cup \Xi}(x) = \max[x_\Psi(x), z_\Xi(x)] \forall x \in \Omega$$
$$z_{\Psi \cap \Xi}(x) = \min[z_\Psi(x), z_\Xi(x)] \forall x \in \Omega$$

A very important concept is the cardinality of a fuzzy set. The cardinality of a fuzzy set, $|\Psi|$, is computed with the same equation as the cardinality of a crisp set:

$$|\Psi| = \sum_{x \in \Omega} z_\Psi(x)$$

The notion of subsethood, the degree to which A is a subset of B, is the fuzzy set approach often proposed for sensor fusion (14). The subsethood theorem states that the degree to which A is a subset of B is the cardinality of A intersected with B divided by the cardinality of A:

$$S(A, B) = \frac{|A \cap B|}{|A|}$$

The power set of a fuzzy set B, P_B, is the set of sets such that

$$A \in P_B \Leftrightarrow z_A(x) \leq z_B(x) \forall x$$

So we can see that if A is an element of the power set of B, the subsethood of A with respect to B will be 1.0.

The subsethood equation above suggests an analogy between subsethood and probability. If we interpret cardinality to be analogous to probability, the numerator of the right hand side is analogous to the joint probability of "A and B" and the denominator to the marginal probability of A; leading to the interpretation that $S(A,B)$is analogous to the conditional probability of Bgiven A, $p(B|A)$. It is easy to show that

$$S(A, B) = \frac{S(B, A)|B|}{|A|}$$

Table 3. Likelihoods for TWS Reports

Focal States	p(rpt 1|state)	p(rpt 2|state)
SAM-X.acq	0.1	0.4
SAM-X.ttr	0.1	0.4
SAM-X.ml	0.1	0.4
SAM-Y.acq	0.6	0.4
SAM-Y.ttr	1.0	0.8
SAM-Y.ml	0.8	0.4

Table 4. Posteriors for First and Both Reports

Focal States	p(state|rpt 1)	p(state|rpts 1 & 2)
SAM-X.acq	0.037	0.027
SAM-X.ttr	0.037	0.027
SAM-X.ml	0.037	0.027
SAM-Y.acq	0.22	0.16
SAM-Y.ttr	0.37	0.54
SAM-Y.ml	0.30	0.22

Table 5. Supportability and Plausability After Both Reports

Focal States	s(state after rpts 1 & 2)	pl(state after rpts 1 & 2)
SAM-X	0.02	0.04
SAM-Y	0.96	0.98
SAM-X.acq	0.00	0.04
SAM-X.ttr	0.00	0.04
SAM-X.ml	0.00	0.04
SAM-Y.acq	0.00	0.29
SAM-Y.ttr	0.49	0.98
SAM-Y.ml	0.00	0.39

furthering the analogy between subsethood and probability because this equation is clearly reminiscent of Bayes's rule.

The following results, in which the fuzzy set E represents evidence and A_i represents the identification states of interest, are also easy to show:

$$0 \le S(E, A_i) \le 1$$
$$S(E, A_i) = 1, \quad \text{if } E \subset A_i (E \in P_{A_i})$$
$$S(E, A_i \cup A_j) = S(E, A_i) + S(E, A_j) - S(E, A_i \cap A_j)$$
$$S(E, A_i \cap A_j) = S(E, A_j) S(A_j \cap E, A_i)$$

These results generalize to multiple items of evidence:

$$S(E_1 \cap E_2, A_i) = \frac{S(E_1, A_i \cap E_2)}{S(E_1, E_2)} = \frac{|E_1 \cap E_2 \cap A_i|}{|E_1 \cap E_2|}$$

Example. For a sample comparison of the Bayesian, evidence theory, and fuzzy set approaches consider the following Threat Warning example that has been used to illustrate evidence theory. There are two types of command guided surface-to-air missile systems about which friendly aircraft are concerned, SAM-X and SAM-Y. The radar of each missile system has three operational states: acquisition (acq), target track (ttr), and missile launch (ml). A radar may be in only one state at a point in time. The friendly aircraft has a threat warning system (TWS) with sensors that monitor such radar parameters as radar frequency (RF) and pulse repetition frequency (PRF) and attempt to determine which SAM radar, and its associated operational state, is painting the aircraft. Suppose that the TWS of the friendly aircraft provides two independent sensor reports (Table 2) about the same SAM site within a relatively short period of time.

There may be several ways to convert this information into likelihoods for application in Bayes's theorem; our approach here is to add the values that do not conflict with the likelihood being calculated. For example,

$$p(\text{report } 1|\text{Y.acq}) = 0.3 + 0.2 + 0.1 = 0.6$$

Table 3 shows the likelihoods for both TWS reports.

Assuming uniform priors over the six focal states, Bayes's theorem yields the results shown in Table 4 under the assumption that we compute a posterior after receiving the first report and then again after both reports.

The evidence theory solution begins by assuming that all prior information is uncommitted. Therefore, the updated uncertainty after the first report will be the first report. The results after the second report are: SAM-X is .02, SAM-Y is .17, SAM-Y.ttr is .49, SAM-Y.acq or SAM-Y.ttr is .10, SAM-Y.ttr or SAM-Y.ml is .20, and uncommitted is .02. The supportability and plausibility after the second report are shown in Table 5.

For the fuzzy solution to this problem we assume that each report element, for example, SAM-Y (any state), is a fuzzy set. Using the extension principle (4), which is a max(min(. . .)) operation, we can compute the fuzzy set associated with each TWS report. Table 6 shows how the fuzzy set for the first report ("Rpt 1") is computed and then shows the results for "Rpt 2" and the combination (intersection) of reports 1 and 2. The fuzzy results, using the subsethood theorem, for report 1 and then for both reports are shown in Table 7.

The results of Bayes's rule, evidence theory, and the subsethood theorem are shown in Table 8. These results are not dissimilar. However, there are situations in which the results will be significantly different.

Table 6. Fuzzy Sets for TWS Reports

	Focal States of Universal Set					
Report Elements & Reports	X.acq	X.ttr	X.ml	Y.acq	Y.ttr	Y.ml
"Y"	0.0	0.0	0.0	0.3	0.3	0.3
"Y.ttr or Y.ml"	0.0	0.0	0.0	0.0	0.4	0.4
"Y.acq or Y.ttr"	0.0	0.0	0.0	0.2	0.2	0.0
"Uncommitted"	0.1	0.1	0.1	0.1	0.1	0.1
"Rpt 1"	0.1	0.1	0.1	0.3	0.4	0.4
"Rpt 2"	0.2	0.2	0.2	0.2	0.4	0.2
"Rpt 1 & Rpt 2"	0.1	0.1	0.1	0.2	0.4	0.2

Table 7. Subsethood for First and Both Reports

Focal States	S(rpt 1, state)	S(rpts 1 & 2, state)
SAM-X.acq	0.072	0.091
SAM-X.ttr	0.072	0.091
SAM-X.ml	0.072	0.091
SAM-Y.acq	0.21	0.18
SAM-Y.ttr	0.29	0.36
SAM-Y.ml	0.29	0.18

Situation and Threat Assessment

Situation and threat refinement are processes of reasoning about aggregations of objects and projecting objects and aggregates of objects forward in time. The locations, identities, activities, and time for which the objects are expected to remain doing those activities are the major characteristics associated with the objects and their aggregates that are of interest.

Reasoning systems are often considered to have three major elements: a knowledge representation scheme, an inference or evaluation process, and a control structure for searching and computation. Table 9 presents many alternate approaches for representing knowledge, inferring or evaluating, and controlling the reasoning process. There is *no* suggested linkage among items on the same row. Rather, this table presents many options for addressing each of the three main elements of a reasoning system. To build a reasoning system, one must select one or more options for representing knowledge, one or more options for conducting inference and evaluation, and one or more methods for controlling the reasoning process.

Table 9. Element Option Table for Reasoning Systems

Knowledge Representation Scheme	Inference or Evaluation Process	Control Structure
Rule	Deduction	Search
Frame	Induction	Reason Maintenance System
Hierarchical Classification	Abduction	Assumption-based Truth Maintenance
Semantic Net	Analogy	Hierarchical Decomposition
Neural Net	Classical Statistics	Control Theory
Nodal Graphs	Bayesian Probability	Opportunistic Reasoning
Options, Goals, Criteria, Constraints	Evidence Theory	Blackboard Architecture
Script	Polya's Plausible Inference	
Time Map	Fuzzy Sets and Fuzzy Logic	
Spatial Relationships	Confidence Factors	
Analytical Model	Decision Theory (Analysis)	
	Circumscription	

ADVANCED APPLICATIONS

While the primary research in data fusion has focused on military applications for detecting and tracking military targets, data fusion processes are being applied in a broad range of civil and commercial applications as well. Image data fusion applications combine multiple images of a common scene or object by registering the imagery to produce enhanced (spatial or spectral) composite imagery, detect changes over time, or to merge multiple video sources. Medical applications include the fusion of magnetic resonance (MR) and computer tomography (CT) images into full 3-D models of a human body for diagnosis and treatment planning. The fusion of geospatial data for mapping, charting, and geodetic applications includes registering and linking imagery, maps, thematic maps, and spatially-encoded text data in a common data base into a geographic information system (GIS). Robotic applications require the registration and combination of imaging, tactile, and other sensors for inspection and manipulation of parts. Financial applications, similar to intelligence uses, require the association of vast amounts of global financial data to model and predict market behaviors for decision analysis.

AREAS OF FURTHER STUDY

Data fusion technology will always be faced with demands to accept higher data rates and volumes as sensing technologies provide higher fidelity data, and data base technologies provide greater capacity to store information. Advanced applications of data fusion will also include integrated sensors, robust processing, learning capabilities, robust spatial data structures, and spatial reasoning. Key areas of further study in data fusion include:

Optimal Sensor and Process Management. The management of complex networks of sensors to achieve optimum information-based performance and operational effectiveness will require advances in the application of optimal search and programming methods. Similarly, as data fusion processing networks grow in complexity, advanced management methods must be developed to allocate diverse networked fusion resources to acquired data sets.

Uncertain Data Management. The ability to quantify uncertainty, combine multiple uncertain data elements,

Table 8. Posterior, Supportability, and Subsethood After Second Report

Focal States	p(state|rpts 1 & 2)	s − pl(state after 2 reports)	S(rpts 1 & 2, state)
SAM-X.acq	0.027	0.00–0.04	0.091
SAM-X.ttr	0.027	0.00–0.04	0.091
SAM-X.ml	0.027	0.00–0.04	0.091
SAM-Y.acq	0.16	0.00–0.29	0.18
SAM-Y.ttr	0.54	0.49–0.98	0.36
SAM-Y.ml	0.22	0.00–0.39	0.18

and infer uncertain information and knowledge requires advances in methods to (1) combine, manage, and represent uncertainty, (2) create and maintain multiple hypotheses, and (3) provide traceability to source data through the use of "pedigree" data.

Dynamic Databases and Information Representation. As the volumes of data to be fused increase, means for mediating between heterogeneous databases must be developed, and more flexible methods of representing information models (text, hypertext, spatial data, imagery, and video, etc.) must be developed.

Knowledge Prediction. High level and adaptive, learning models of complex processes must be developed to predict the behavior of groups and complex relationships beyond the level of simple level 1 target tracks.

Visualization. The results of most data fusion systems must ultimately be presented to human decision-makers, requiring advances in the methods to efficiently display high volumes of complex information and derived knowledge, and to provide the ability to "drill-down" to the underlying data fusion processes and sources of data.

BIBLIOGRAPHY

1. E. L. Waltz and J. Llinas, *Data Fusion with Military Applications,* Norwood, MA: Artech House, 1990.
2. D. L. Hall, *Mathematical Techniques in Multisensor Data Fusion,* Norwood, MA: Artech House, 1992.
3. B. V. Dasarathy, *Decision Fusion,* Washington, DC: IEEE Computer Society Press, 1994.
4. R. C. Luo and M. G. Kay, *Multisensor Integration and Fusion for Intelligent Machines and Systems,* Norwood, NJ: Ablex, 1995.
5. D. L. Hall and J. Llinas, An Introduction to Multisensor Data Fusion, *Proc. IEEE,* **85**: 6–23, 1997.
6. S. S. Blackman, *Multiple-Target Tracking with Radar Applications,* Norwood, MA: Artech House, 1986.
7. Y. Bar-Shalom and X. Li, *Multitarget Multisensor Tracking Principles and Techniques,* Storrs, CT: YBS Publishing, 1995.
8. K. R. Pattipati, et al., A New Relaxation Algorithm and Passive Sensor Data Association, *IEEE Trans. Autom. Control,* **37**: 198–213, 1992.
9. D. M. Buede and P. Girardi, A Target Identification Comparison of Bayesian and Dempster-Shafer Multisensor Fusion, *IEEE Trans. Syst. Man Cybern,* **27A**: 569–577, 1997.
10. J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference,* San Mateo, CA: Morgan Kaufmann, 1988.
11. A. P. Dempster, A Generalization of Bayesian Inference, *J. Roy. Statistical Soc.,* Series B, **30** (2): 205–247, 1968.
12. G. Shafer, *A Mathematical Theory of Evidence,* Princeton, NJ: Princeton Univ. Press, 1976.
13. L. A. Zadeh, Fuzzy Sets, *Inf. Control,* **8**: 338–353, 1965.
14. B. Kosko, *Neural Networks and Fuzzy Systems,* Englewood Cliffs, NJ: Prentice-Hall, 1992.

DENNIS M. BUEDE
George Mason University

EDWARD L. WALTZ
Environmental Research Institute of Michigan

DATA MANAGEMENT. See DATABASES.

DATA MART. See DATAWAREHOUSING.

DATA MINING. See DATA REDUCTION; MACHINE LEARNING.

DATA MODELS, OBJECT-ORIENTED. See OBJECT-ORIENTED DATABASES.