6

Data Fusion Automation: A Top-Down Perspective

6.1 Introduction 6-1
Biological Fusion Metaphor • Command and Control Metaphor • Puzzle-Solving Metaphor • Evidence Combination • Information Requirements • Problem Dimensionality • Commensurate and Noncommensurate Data
6.2 Biologically Motivated Fusion Process Model 6-8
6.3 Fusion Process Model Extensions 6-14
Short-, Medium-, and Long-Term Knowledge • Fusion Classes • Fusion Classes and Canonical Problem-Solving Forms
6.4 Observations 6-22
Observation 1 • Observation 2 • Observation 3 • Observation 4 • Observation 5
Acknowledgment 6-25
References 6-25

Richard Antony
VGS Inc.

6.1 Introduction

This chapter offers a conceptual-level view of the data fusion process and discusses key principles associated with both data analysis and information combination. The discussion begins with a high-level view of data fusion requirements and analysis options. Although the discussion focuses on tactical situation awareness development, a much wider range of applications exists for this technology.

After motivating the concepts behind effective information combination and decision making through a series of easily understood metaphors, the chapter

- Presents a top-down view of the data fusion process,
- Discusses the inherent complexities of combining uncertain, erroneous, and fragmentary information,
- Offers a taxonomic approach for distinguishing classes of fusion algorithms, and
- Identifies key algorithm requirements for practical and effective machine-based reasoning.

6.1.1 Biological Fusion Metaphor

Multiple sensory fusion in biological systems provides a natural metaphor for studying artificial data fusion systems. As with any good metaphor, consideration of a simpler or more familiar phenomenon can provide valuable insight into the study of a more complex or less familiar process.

0-8493-2379-7/01/$0.00+$1.50

Even the most primitive animals sense their environment, develop some level of situation awareness, and react to the acquired information. Situation awareness directly supports survival of the species by assisting in the acquisition of food and the avoidance of animals of prey. A barn owl, for instance, fuses visual and auditory information to help accurately locate mice under very low light conditions, while a mouse responds to threatening visual and auditory cues to attempt to avoid being caught by an owl.

In general, natural selection has tended to favor the development of more capable senses (sensors) and more effective utilization of the derived information (exploitation and fusion). Color vision in humans, for instance, is believed to have been a natural adaptation that permitted apes to more easily locate ripe fruit among vegetation. Situation awareness in animals can rely on a single, highly developed sense, or on multiple, often less capable senses. A hawk depends principally on a highly acute visual search and tracking capability, while a shark primarily relies on its sense of smell when hunting. Sexual attraction can depend primarily on sight (plumage), smell (pheromones), or sound (mating call). For humans, sight is arguably the most vital sense, with hearing a close second. Dogs, on the other hand, rely most heavily on the senses of smell and hearing, with vision typically acting as a secondary information source.

Sensory input in biological organisms typically supports both sensory cueing and situation awareness development. Sounds cue the visual sense to the presence and the general direction of an important event. Information gained by the aural sense (i.e., direction, speed, and tentative object classification) is then combined (fused) with the information gathered by the visual system to produce more complete, higher confidence, or higher level situation awareness. In many cases, multiple sensory fusion can be critical to successful decision making. Food that looks appetizing (sight) might be extremely salty (taste), spoiled (smell), or too hot (touch). At the other extreme, fusion of multiple sensory input might be unnecessary if the various senses provide highly redundant information. Bacon frying in a pan need not be seen, smelled, and tasted to be positively identified; each sense, taken separately, could perform such a function.

Although discarding apparently redundant information may seem to be prudent, such information can aid in sorting out conflicts, both intentional (deception) and unintentional (confusion). While single-source deception is reasonably straightforward to perpetrate, deception across multiple senses (sensor modalities) is considerably more difficult. For example, successful hunting and fishing depend, to a large degree, on effective multisource deception. Duck hunters use both visual decoys and mating calls to simultaneously provide deceptive visual and auditory information. Because deer can sense danger through the sense of smell, sound, and sight, the shrewd hunter must mask his scent (or stay down-wind), make little or no noise, and remain motionless if the deer looks in his direction. Even in nonadversarial applications, data fusion requires resolution of unintentional conflicts among supporting data sources in order to deal effectively with the inherent uncertainty in both the measurement and decision spaces.

Multiple sensory fusion need not be restricted to the familiar five senses of sight, sound, smell, taste, and touch. Internal signals, such as acidity of the stomach, coupled with visual and/or olfactory cues, can trigger hunger pains. The fusion of vision, inner-ear balance information, and muscle feedback signals facilitate motor control. In a similar manner, measurement and signature intelligence (MASINT) in a tactical application focuses on the collection and analysis of a wide range of nontraditional information classes.

6.1.2 Command and Control Metaphor

The game of chess provides a literal metaphor for military command and control (C^2), as well as an abstract metaphor for any system that senses and reacts to its environment. Both chess players and battlefield commanders require a clear picture of the "playing field" to properly evaluate the options available to them and their opponents. In both chess and C^2, opposing players command numerous individual resources (i.e., pieces or units) that possess a range of characteristics and capabilities. Resources and strategies vary over time. Groups of chess pieces are analogous to higher-level organizations on the battlefield. The chessboard represents domain constraints to movement that are similar to constraints posed by terrain, weather, logistics, and other features of the military problem domain. Player-specific strategies are analogous to tactics, while legal moves represent established doctrine. In both domains, the overall objective of an opponent may be known, while specific tactics and subgoals must be deduced.

Despite a chess player's complete knowledge of the chess board (*all domain constraints*), the location of all pieces (*own and opponent-force locations*), and all legal moves (*own and opponent-force doctrine*), and his ability to exercise direct control over all of his own assets, chess remains a highly challenging game. Metaphorically similar to chess, tactical situation development has numerous domain characteristics that make it an even more challenging problem.

First, battlefield commanders normally possess neither a complete nor fully accurate picture of their own forces or those of their adversaries. Forced to deal with incomplete and inaccurate force structure knowledge, as well as location uncertainty, chess players would be reduced to guessing the location and composition of an adversary's pieces, somewhat akin to playing "Battleship," the popular children's game.

Second, individual sensors provide only limited observables, coverage, resolution, and accuracy. Thus, the analysis of individual sensor reports tend to lead to ambiguous and rather local interpretations. Third, domain constraints in tactical situation awareness are considerably more complex than the well-structured (and level) playing field in chess. Fourth, doctrinal knowledge in the tactical domain tends to be more difficult to exploit effectively and far less reliable that its counterpart in chess.

A wide range of other application-motivated metaphors can also be useful for studying specific fusion applications. Data fusion, for example, seems destined to play a significant role in the development of future "smart highway" control systems where a simple car driving metaphor can be applied to study sensor requirements and fusion opportunities. The underpinning of such a system is a sophisticated control capability that optimally resolves a range of conflicting requirements, such as (1) expedite the movement of both local and long distance traffic, (2) ensure maximum safety for all vehicles, and (3) create the minimum environmental impact. The actors in the metaphor are drivers (or automated vehicle control systems), the rules of the game are the "rules of the road," and domain constraints are the road network and traffic control means. Individual players possess individualized objectives and tactics; road characteristics and vehicle performance capabilities provide physical constraints on the problem solution.

6.1.3 Puzzle-Solving Metaphor

Situation awareness development requires the production and maintenance of an adequate multiple level-of-abstraction picture of a (dynamic) situation; therefore, the data fusion process can be compared to assembling a complex jigsaw puzzle for which no picture of the completed scene exists. While assembling puzzles that contain hundreds of pieces (*information fragments*) can challenge an individual's skill and patience, the production of a comprehensive situational picture, created by fusing disparate and fragmentary sensor-derived information, represents an even more challenging task. Although a completed jigsaw puzzle represents a fixed scene, the process of collecting and integrating the numerous information fragments clearly evolves over time. Time, on the other hand, represents a key dimension in highly dynamic tactical situation awareness applications.

The partially completed puzzle (*fused situation awareness product*) illustrated in Figure 6.1 contains numerous aggregate objects (i.e., forest and meadow), each composed of simpler objects (i.e., trees and ground cover). Each of these objects, in turn, have been assembled from multiple puzzle pieces, some representing a section of bark on a single tree trunk, others a grassy area associated with a meadow. In terms of the metaphor then, sensor-derived information can be associated with individual puzzle pieces, providing little more information than color and texture, as well as pieces that depict higher level of abstraction objects.

At the beginning of the reconstruction process, problem solving necessarily relies on general analysis strategies (e.g., locate border pieces). Because little context exists to direct either puzzle piece selection or puzzle piece placement, at the early stages of the process, rather simple, brute-force pattern matching strategies are needed. A predominately blue-colored piece, for example, might represent either sky or water with little basis for distinguishing between the two interpretations. Unless they came from an unopened box, there may be no assurance that the scattered pieces on the table all belong in the puzzle

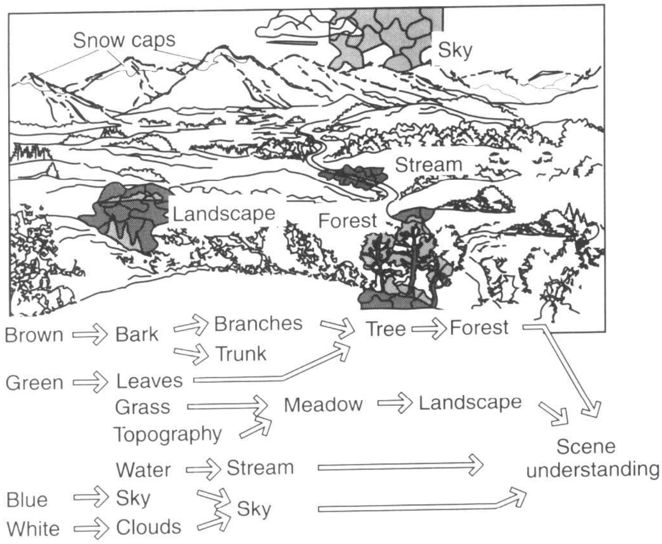

FIGURE 6.1 Puzzle-solving metaphor example.

under construction. However, once certain sections of the puzzle have been filled in, the assembly process (*fusion*) tends to become much more goal-directed.

Fitting a single puzzle piece supports both scene entropy reduction as well as higher level-of-abstraction scene interpretation. As regions of the puzzle begin to take form, identifiable features in the scene emerge (e.g., trees, grass, and cliffs) and higher-level interpretations can be developed (e.g., forest, meadows, and mountains). By supporting the placement of the individual pieces, as well as the goal-driven search (*sensor resource management*) for specific pieces, the *context* provided by the developing multiple level-of-abstraction picture of the scene (*situation awareness product*) helps further focus the reconstruction process (*fusion process optimization*).

Just as duplicate or erroneous pieces can significantly complicate puzzle assembly, redundant and irrelevant sensor-derived information similarly burdens machine-based situation development. Therefore, goal-directed information collection offers a two-fold benefit: critical information requirements are satisfied and the collection (and subsequent analysis) of unnecessary information is minimized. Although numerous puzzle pieces may be yet unplaced (*undetected objects*) and perhaps some pieces are actually missing (*information not collectible by the available sensor suite*), a reasonably comprehensive, multiple level-of-abstraction understanding of the overall scene (*situation awareness*) gradually emerges.

Three broad classes of knowledge are apparent in the puzzle reconstruction metaphor:

- Individual puzzle pieces — collected information fragments, i.e., *sensor-derived knowledge*,
- Puzzle-solving strategies, such as edge detection and pattern matching — a priori *reasoning knowledge*
- World knowledge, such as the relationship between meadows and grass — *domain context knowledge*.

To investigate the critical role that each knowledge form plays in fusion product development, recast the analysis in terms of a building construction metaphor. Puzzle pieces (*sensor input*) are clearly the building blocks required to assemble the scene (*fused situation awareness product*). *A priori* reasoning knowledge represents construction knowledge and skills, and context provides the nails and mortar that "glue" the sensor input together to form a coherent whole. When too many puzzle pieces (or building blocks) are missing (*inadequate sensor-derived information*), scene reconstruction (or building construction) becomes difficult or impossible.

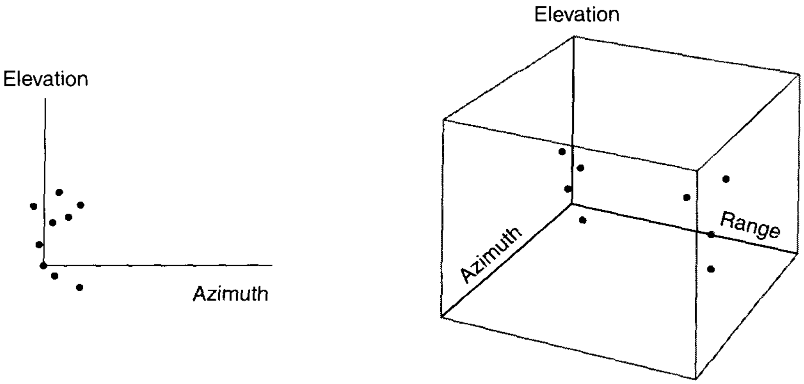

FIGURE 6.2 (a) Two-dimensional measurements and (b) the corresponding three-dimensional measurement space.

A simple example demonstrates how both the complexity of the fusion process and the quality of the resultant product are sensitive to the availability of adequate information. Figure 6.2(a) illustrates a cluster of azimuth and elevation measurements associated with two separate groups of air targets. Given the spatial overlap between the data sets, reliable target-to-group assignment may not be possible, regardless of the selected analysis paradigm or the extent of algorithm training. However, with the addition of range measurements (*increased measurement space dimensionality*), two easily separable clusters become readily apparent (Figure 6.2(b)). Because the information content of the original 2-D data set was fundamentally inadequate, even sophisticated clustering algorithms would be unable to discriminate between the two target groups. However, with the addition of the third measurement dimension, a simple clustering algorithm easily handles the decision task.

Reasoning knowledge can be implemented using a spectrum of problem solving paradigms (e.g., rules, procedures, and statistical-based algorithms), evidence combination strategies (e.g., Bayes, Dempster-Shafer, and fuzzy set theory), and decision-making approaches (e.g., rule instantiation and parametric algorithms). In general, the process of solving a complex puzzle (or performing automated situation awareness) benefits from **both** bottom-up (deductive-based) and top-down (goal-directed) reasoning that exploits relationships among the hierarchy of domain entities (i.e., primitive, composite, aggregate, and organizational).

In the puzzle-solving metaphor, *context knowledge* refers to relevant domain knowledge **not** explicitly contained within a puzzle piece (*non-sensor-derived knowledge*). Humans routinely apply a wide range of contextual knowledge during analysis and decision making.* For example, context-sensitive evaluation of Figure 6.1 permits the determination that the picture is a summer scene in the western U.S. The season and location are deduced from the presence of deciduous trees in full leaf (summer) in the foreground and jagged snow-capped mountain peaks in the distance (western U.S.). In a similar fashion, the exploitation of context knowledge in automated fusion systems can promote much more effective and comprehensive interpretations of sensor-derived information.

In both puzzle assembly and automated situation development, determining when an *adequate* situation representation has been achieved can be difficult. In the puzzle reconstruction problem, although the general landscape characteristics might be evident, missing puzzle pieces could depict denizens of the woodland community that can be hypothesized, but for which no compelling evidence yet exists. On the other hand, individual puzzle pieces might contain partial or ambiguous information. For

---

* This fact partially accounts for the disparity in performance between manual and automated approaches to data fusion.

example, the presence of a section of log wall in the evolving scene suggests the possibility of a log cabin. However, additional evidence is required to validate such a hypothesis.

6.1.4 Evidence Combination

Reliance on a single information source can lead to ambiguous, uncertain, and inaccurate situation awareness. Data fusion seeks to overcome such limitations by synergistically combining all relevant (and available) information sources leading to the generation of consistent, accurate, comprehensive, and global situation awareness. A famous poem by John Godfrey Saxe,* written more than a century ago, aptly demonstrates both the need for and challenge of effectively combining fragmentary information.

The poem describes an attempt by six blind men to gain a first-hand understanding of an elephant. The first man happens to approach the elephant from the side and surmises that an elephant must be something like a wall. The second man touches the tusk and imagines an elephant to be like a spear. The third man approaches the trunk and decides an elephant is similar to a snake. The fourth man reaches out and touches a leg and determines an elephant to be much like a tree. The fifth man chances to touch an ear and imagines an elephant must be like a fan. The sixth man grabs the tail and concludes an elephant is similar to a rope. While each man's assessment is entirely consistent within his own limited sensory space and myopic frame of reference, unless the six observations are effectively integrated (*fused*), a true picture of an elephant fails to emerge.

Among other insights, the puzzle-solving metaphor illustrated that (1) complex dependencies can exist among and between information fragments and the completed situation description, and (2) determining whether an individual puzzle piece actually belongs to the scene being assembled can be difficult. Even when the collected information is known to be relevant, based strictly on local interpretations, determining whether a given blue-colored piece represents sky, water, or some other feature class may not be possible. Much like assembling observations, hunting for clues, and evaluating motives required during criminal investigations, a similar approach to information combination is required by general situation awareness systems. Just as at the outset of a criminal investigation, a single strand of hair might appear insignificant, but it could later prove to be the key piece of evidence that discriminates among several suspects. Similarly, a seemingly irrelevant piece of sensor-derived information might ultimately link observations with motives, or provide other significant situational awareness benefits. Thus, not only is the information content (*information measure*) associated with a given piece of data important; its relationship to the overall fusion task is also vital to achieving successful information fusion. As a direct consequence of this observation, the development of a comprehensive information theoretical framework for the data fusion process appears to be problematic. Only through a top-down, holistic treatment of the analysis task can the content of a single information fragment be properly assessed and its true value to the overall fusion process be fully realized.

6.1.5 Information Requirements

Because no widely accepted formal theory exists for determining when adequate information has been assembled to support a given fusion task, empirical measures of performance generally must be relied upon to evaluate the effectiveness of both individual fusion algorithms and an overall fusion system. In general, data fusion performance can be enhanced by

- Technical improvements in sensor measurements (i.e., longer range, higher resolution, improved signal-to-noise ratio, better accuracy, higher reliability);
- Increased measurement space dimensionality afforded by heterogeneous sensors that provide at least partially independent information;

---

* Saxe, J. G., "The Blind Man and the Elephant," *The Poetical Works of John Godfrey Saxe*, Boston, MA: Houghton, Mifflin and Company, 1882.

- Spatially distributed sensors providing improved coverage, perspective, and measurement reliability;
- Relevant non-sensor-derived domain knowledge to constrain the information combination and decision-making process.

In general, effective data fusion automation requires the development of robust, context-sensitive algorithms that are practical to implement. The first two requirements reflect the "quality of performance" of the algorithm, while the latter reflects cost/benefit tradeoffs associated with meeting a wide range of implicit and explicit performance objectives. In general, *robust* performance argues for the use of all potentially relevant sensor-derived information sources and reasoning knowledge. Achieving *context-sensitive* performance argues for maximal utilization of relevant non-sensor-derived information. On the other hand, to be *practical* to implement and efficient enough to employ in an operational setting, the algorithms may need to compromise some fusion performance quality. Consequently, system developers must quantify or otherwise assess the value of these various information sources in light of system requirements, moderated by programmatic, budgetary, and performance constraints (e.g., decision timeline and hardware capability). The interplay between achieving optimal algorithm robustness and context-sensitivity, on the one hand, and a practical implementation, on the other, is a fundamental tension associated with virtually any form of machine-based reasoning directed at solving complex, real-world problems.

6.1.6 Problem Dimensionality

Effective situational awareness, with or without intentional deception, generally benefits from the collection and analysis of a wide range of observables. As a result of the dynamic nature of many problem domains, observables can change with time and, in some cases, may require continuous monitoring. In a tactical application, objects of interest can be stationary (fixed or currently nonmoving), quasistationary (highly localized motion), or moving. Individual objects possess characteristics that constrain their behavior. Objects emit different forms of electromagnetic energy that vary with time and can indicate the state of the object. Object emissions include *intentional* or active emissions, such as radar, communications, and data link signals, as well as *unintentional* or passive emissions, such as acoustic, magnetic, or thermal signatures generated by internal heat sources or environmental loading. Patterns of physical objects and their behavior provide indications of organization, tactics, and intent. Patterns of emissions, both active and passive, can reveal the same. For example, a sequence of signals emitted from a surface-to-air missile radar over time representing search, lock-on, launch, and hand-over clearly indicates hostile intent.

A single sensor modality is incapable of measuring all relevant information dimensions; therefore, multiple sensor classes often must be relied upon to detect, track, classify, and infer the likely intent of a host of objects, from submarines and surface vessels, to land, air, and space-based objects. Certain sensor classes lend themselves to surveillance applications, providing both wide-area and long-range coverage, and readily automated target detection capability. Examples of such sensor classes include signals intelligence (SIGINT) for collecting active emissions, moving target indication (MTI) radar for detecting and tracking moving targets against a high clutter background, and synthetic aperture radar (SAR) for detecting stationary targets. Appropriately cued, other sensor classes that possess narrower fields of view and that typically operate at much shorter ranges may be capable of providing higher fidelity measurement to support refined analysis. Geospatial and other intelligence databases can provide the static domain context within which the target-sensed data must be interpreted, while environmental sensors generate dynamic context estimates, such as weather and current atmospheric conditions.

6.1.7 Commensurate and Noncommensurate Data

Although the fusion of *similar* (commensurate) information would seem to be more straightforward than the fusion of *dissimilar* (noncommensurate) information, that is not always the case. Three examples are offered to highlight the varying degrees of difficulty associated with the combination of multiple-source data. First, consider the relative simplicity of fusing registered electronic intelligence (ELINT) data and real-time synthetic aperture radar (SAR) imagery. Although these sensors measure dramatically different information dimensions, both sources provide reasonably wide area coverage, relatively good geolocation, and highly complementary information. As a consequence, the fusion process tends to be straightforward. Even when an ELINT sensor provides little more than target line-of-bearing, the ELINT and SAR measurements can potentially be combined by simply overlaying the two data sets. If the line-of-bearing intercepts a single piece of equipment in the SAR image, the radar system class, as well as its precise location, would be known. This information, in turn, can support the identification of other nearby objects in the image (e.g., missile launchers normally associated with track-while-scan radar).

At the other end of the spectrum, the fusion of information from two or more identical sensors can present a significant challenge. Consider, for example, fusing data sets obtained from spatially separated forward-looking infrared (FLIR) radars. Although FLIR imagery provides good azimuth and elevation resolution, it does not directly measure range. Because the range and view angles to targets will be different for multiple sensors, combining such data sets demands sophisticated registration and normalization.

Finally, consider the fusion of two bore-sited sensors: light-intensified and forward-looking infrared (FLIR). The former device amplifies low intensity optical images to enhance night vision. When coupled with the human's natural ability to separate moving objects from the relatively stationary background, such devices permit visualization of the environment and detection of both stationary and moving objects. However, such devices offer limited capability for the detection of stationary personnel and equipment located in deep shadows or under extremely low ambient light levels (e.g., heavy cloud cover, no moon, or inside buildings). Rather than detecting reflected energy, FLIR devices detect thermal radiation from objects. Consequently, these devices support the detection of humans, vehicles, and operating equipment based on their higher temperature relative to the background. Consequently, with bore-sighted sensors, pixel-by-pixel combination of the two separate images may be feasible, providing a highly effective night vision capability.

6.2 Biologically Motivated Fusion Process Model

A hierarchically organized *functional-level* model of data fusion is presented in Chapter 2. In contrast, this section focuses on a *process-level* model. While the functional model describes **what** analysis functions or processes need to be performed, a process-level model describes at a high level of abstraction **how** this analysis is accomplished.

The goal of data fusion, as well as most other forms of data processing, is to turn data into useful information. In perhaps the simplest possible view, all of the required information is assumed to be present within a set of sensor measurements. Thus, the role of data fusion is extraction of information embedded in a data set (separating the wheat from the chaff). In this case, fusion algorithms can be characterized as a function of

- Observables
- Current situation description (e.g., target track files and current situation description)
- *A priori* declarative knowledge (e.g., distribution functions, templates, constraint sets, filters, and decision threshold values).

As shown in Figure 6.3(a), the fusion process output provides updates to the *situation description*, as well as feedback to the *reasoning knowledge base* to support knowledge refinement (learning).

Signal processing, statistical hypothesis testing, target localization performed by intersecting two independently derived error ellipses, and target identification based on correlation of an image with a set of rigid templates are simple examples of such a fusion model. In general, this "information extraction" view of data fusion makes a number of unstated, simplifying assumptions including the existence of

- Adequate information content in the sensor observables
- Adequate sensor update rates
- Homogeneous sensor data

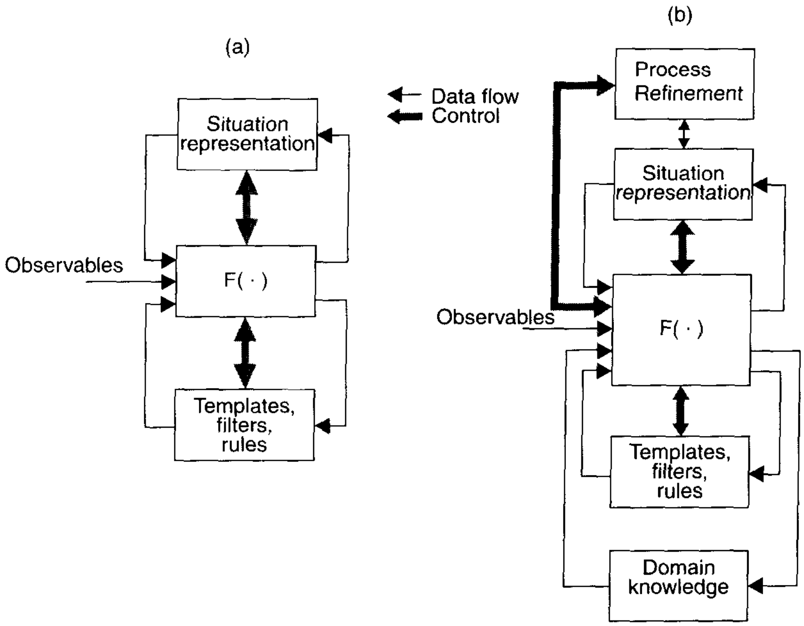

FIGURE 6.3 (a) Basic fusion process model and (b) generalized process model.

- Relatively small number of readily distinguishable targets
- Relatively high resolution sensors
- High reliability sensors
- Full sensor coverage of the area of interest
- Stationary, Gaussian random interference.

When such assumptions are appropriate, data analysis tends to be straightforward and an "information extraction" fusion model is adequate. Rigid template-match paradigms typically perform well when a set of observables closely matches a single template and are uncorrelated with the balance of the templates. Track association algorithms perform well against a small number of moving, widely spaced targets provided the radar generates relatively high update rates. The combination of similar features is often more straightforward than the combination of disparate features. When the sensor data possesses adequate information content, high confidence analysis is possible. High signal-to-noise ratios tend to enhance signal detection. High resolution sensors reduce ambiguity and uncertainty with respect to feature measurements (e.g., location and frequency). High reliability sensors maximize sensor availability. Adequate sensor coverage provides a "complete" view of the areas of interest. Statistical-based reasoning is generally simplified when signal interference can be modeled as a Gaussian random process.

Typical applications where such assumptions are realistic, include

- Track assignment in low target-density environments or for ballistic targets that obey well-established physical laws of motion
- Classification of military organizations based on associated radio types
- Detection of signals and targets exhibiting high signal-to-background ratio.

However, numerous real-world data fusion tasks exhibit one or more of the following complexities:

- Large number of target and nontarget entities (e.g., garbage trucks may be nearly indistinguishable from armored personnel carriers);
- Within-class variability of individual targets (e.g., hatch open vs. hatch closed);

- Low data rates (exacerbating track association problems);
- Multiple sensor classes (disparate numeric and symbolic observables can be difficult to combine);
- Inadequate sensor coverage of areas of interest (i.e., inadequate number of sensors, obscuration due to terrain and foliage, radio frequency interference, weather, or counter measures);
- Inadequate set of sensor observables (e.g., inadequate input space dimensionality);
- Inadequate sensor resolution;
- Registration and measurement errors;
- Inadequate *a priori* statistical knowledge (e.g., unknown prior and conditional probabilities, multimodal density functions, or non-Gaussian and nonstationary statistics);
- Processing and communication latencies;
- High level-of-abstraction analysis product required (i.e., not merely platform location and identification);
- Complex propagation phenomenon (i.e., multipath, diffraction, or atmospheric attenuation);
- Purposefully deceptive behavior.

When such complexities exist, sensor-derived information tends to be incomplete, ambiguous, erroneous, and difficult to combine and/or abstract. Thus, a data fusion process that relies on rigid composition among (1) the observables, (2) the current situation description, and (3) a set of rigid templates or filters, tends to be fundamentally inadequate.

As stated earlier, rather than simply "extracting" information from sensor-derived data, effective data fusion requires the combination, consolidation, organization, and abstraction of information. Such analysis can enhance the fusion product, its confidence, and its ultimate utility in at least four ways:

1. Existing sensors can be improved to provide better resolution, accuracy, sensitivity, and reliability.
2. Additional similar sensors can be employed to improve the coverage and/or confidence in the domain observables.
3. Dissimilar sensors can be used to increase the dimensionality of the observation space, permitting the measurement of at least partially independent target attributes (a radar can offer excellent range and azimuth resolution, while an ELINT sensor can provide target identification).
4. Additional domain knowledge and context constraints can be utilized.

While the first three recommendations effectively *increase* the information content and/or dimensionality of the observables, the latter effectively *reduces* the decision space dimensionality by constraining the possible decision states.

Observables can be treated as *explicit* knowledge (i.e., knowledge that is explicitly provided by the sensors). Context knowledge, on the other hand, represents *implicit* (or non-sensor-derived) knowledge. Although human analysts routinely use both forms in performing fusion tasks, automated approaches have traditionally relied almost exclusively on the former.

As an example of the utility of implicit domain knowledge, consider the extrapolation of the track of a ground-based vehicle that has been observed moving along the relatively straight-line path shown in Figure 6.4. Although the target is a wheeled vehicle traveling along a road with a hairpin curve just beyond the last detection point, a purely statistical-based tracker will likely attempt to extend the track through the hill (the reason for the curve in the road) and into the lake on the other side.

Although tracking aircraft, ballistic projectiles, and naval vessels using statistical-based motion models has been highly successful, adapting such algorithms to tracking ground vehicles has proved to be a considerable challenge. Tracked and wheeled vehicles typically exhibit many more degrees of freedom than a high performance aircraft or naval vessel because they can stop and move in an unpredictable manner. Additional complications include the potentially large numbers of ground vehicles, nonresolvable individual vehicles, terrain and vegetation masking, and infrequent target update rates. However, through the application of relevant domain constraints (e.g., mobility, observability, vehicle class behavior, and vehicle group behavior), the expectation-based analysis process can be effectively constrained,

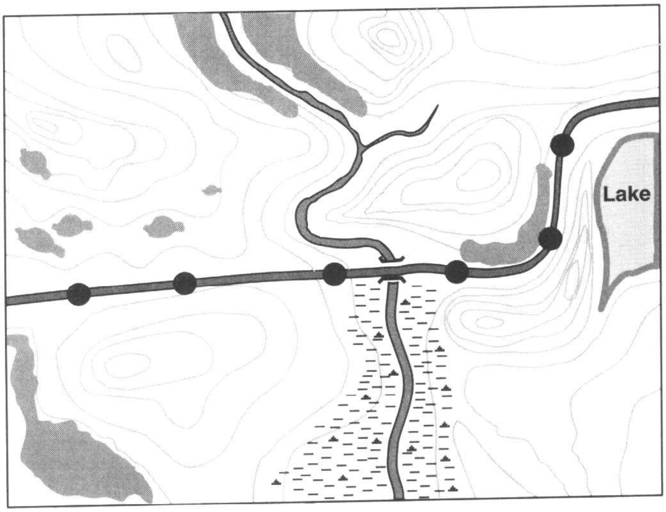

FIGURE 6.4 Road-following target tracking model.

thus helping to manage the additional degrees of freedom. In much the same way that a system of equations with too many unknowns does not produce a unique solution, "missing" domain knowledge can lead to an "underdamped" Kalman filter solution to ground target tracking. In recognition of the benefits of context-sensitive analysis, domain-sensitive ground target tracking models have received considerable interest in recent years.

In addition to the importance of reasoning in context, the road-following target tracking problem also dramatically illustrates the critical role of *paradigm selection* in the algorithm development process. Rather than demonstrating the failure of a statistical-based tracker, the above example illustrates its misapplication. Applying a purely statistical approach to this problem assumes (perhaps unwittingly) that domain constraints are either irrelevant or insignificant. However, in this application, domain constraints tend to be stronger than the relatively weak constraints on platform motion provided by a strictly statistical-based motion model.

Paradigm selection, in fact, must be viewed as a key component of successful data fusion automation. Consequently, algorithm developers must ensure that both the capability and limitations of a selected problem-solving paradigm are appropriately matched to the requirements of the fusion task they are attempting to automate.

To illustrate the importance of both context-sensitive reasoning and paradigm selection, consider the problem of analyzing the time-stamped radar detections from multiple closely spaced targets, some with potentially crossing trajectories, as illustrated in Figure 6.5. A traditional statistical tracking algorithm typically associates the "closest" (with respect to a specified evaluation metric) new detection to an existing track. A human analyst, on the other hand, would quite naturally invoke a context-sensitive model of vehicle behavior. By employing multiple behavior models, alternative interpretations of the observations can be made. False hypotheses can be eliminated once adequate information is obtained to resolve the associated ambiguity.

Emulating such an analysis strategy requires the time-stamped detections to be associated with local cultural and topographic features. In addition, the analysis model(s) must accommodate individual vehicle-class capabilities, as well as *a priori* class-specific behavioral knowledge. By doing so, it can be inferred that tracks 1–3 would be highly consistent with a road-following behavior, tracks 4 and 5 would be determined to be most consistent with a minimum terrain-gradient following behavior, while track 6 would be found to be inconsistent with any ground-based vehicle behavior model. By evaluating track updates from targets 1–3 with respect to road association, estimated vehicle speed, and observed inter-target spacing (assuming individual targets are resolvable), it can be deduced that targets 1–3 are wheeled

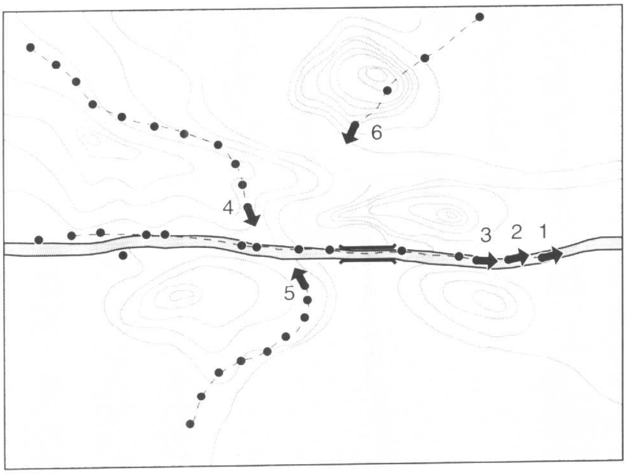

FIGURE 6.5 Example of the fusion of multiple-target tracks over time.

vehicles traveling in a convoy along a secondary road. Based on the maximum observed vehicle speeds and the associated surface conditions along their trajectories, tracks 4 and 5 can be deduced to be tracked vehicles. Finally, because of its relatively high speed and the rugged terrain in the vicinity, track 6 would be determined to be most consistent with a low-flying airborne target. Because the velocity of target 6 is too low to be a fixed-wing aircraft, the target can be inferred to be a helicopter.

Targets may be moving at one instant of time and stationary at another and communicating during one interval and silent during another, resulting in four mutually exclusive target states: (1) moving, nonemitting, (2) moving emitting, (3) nonmoving, nonemitting, and (4) nonmoving, emitting. Over time, many entities in the domain may change between two or more of these four states. Thus, if the situation awareness product is to be continuously maintained, data fusion inherently involves a recursive analysis. Table 6.1 provides a mapping between these four target states and a wide range of sensor classes. As shown, the ability to track entities through these state changes effectively requires multiple source sensor data.

In general, individual targets exhibit complex patterns of behavior that can help discriminate object classes and identify activities of interest. Consider the scenario depicted in Figure 6.6, showing the movement of a tactical erectable missile launcher (TEL) between time t_0 and time t_6. At t_0, the vehicle is in a location that makes it difficult to detect. At t_1, the vehicle is moving along a dirt road at velocity v_1. At time t_2, the vehicle continues along the road and begins communicating with its support elements. At time t_3, the vehicle is traveling off road at velocity v_3 along a minimum terrain gradient path. At time t_4, the target has stopped moving and begins to erect its launcher. At time t_5, just prior to launch, radar emissions begin. At time t_6, the vehicle is traveling to a new hide location at velocity v_6.

TABLE 6.1 Mapping between Sensor Classes and Target States

	Sensor Classes								
Target Classes	MTI Radar	SAR	Laser Radar	COMINT	ELINT	FLIR	Optical	Acoustic	Measurement and Signature
Moving/emitting	•		•	•	•	•		•	•
Moving/nonemitting	•		•			•		•	
Nonmoving/emitting		•		•	•	•	•		•
Nonmoving/nonemitting		•				•	•		

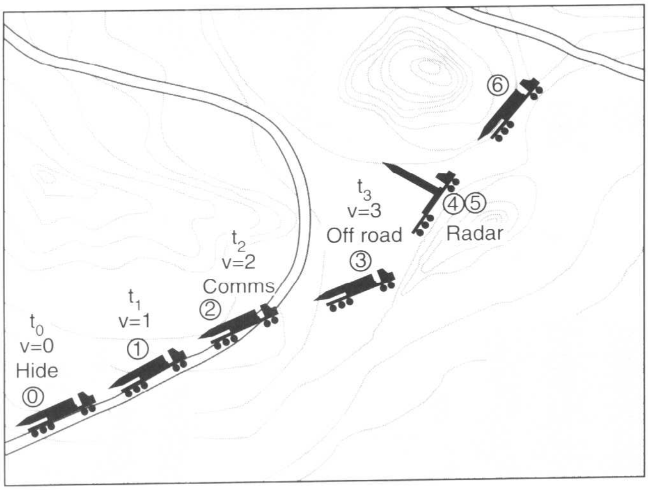

FIGURE 6.6 Dynamic target scenario showing sensor snapshots over time.

TABLE 6.2 Interpretation of Scenario Depicted in Figure 6.6

1	2	3	4	5	6	7	8
State	State Class	Velocity	Emission	Potentially Contributing Sensors	Local Interpretation	High-level Interpretation	Global Interpretation
0	Nonmoving/ nonemitting	0		SAR FLIR Imagery Video	Light foliage	Concealment	Hide
1	Moving/ nonemitting	v_1		MTI FLIR Video	Road association	High-speed mobility	Move to launch
2	Moving/ emitting	v_2	Comm type 1	MTI FLIR Video COMINT	Road association C2 network active	High-speed mobility Coordination status	
3	Moving/ nonemitting	v_3		MTI FLIR Video	Off road Good mobility	Minimum terrain gradient path Local goal seeking	
4	Nonmoving/ nonemitting	v_4		MTI FLIR Imagery	Open, flat Good mobility Good visibility	Tactical activity or staging area	Launch preparation and launch
5	Nonmoving/ emitting	v_5	Comm type 1 and 2; radar	SAR FLIR Imagery Video SIGINT	Coordination status Prelaunch transmission	Launch indication	
6	Moving/ nonemitting	v_6		MTI FLIR Video	High-speed travel	Rapid movement Road seeking	Move to hide

Table 6.2 identifies sensor classes that could contribute to the detection and identification of the various target states. Opportunities for effective sensor cross cueing for the TEL scenario discussed earlier are shown in the "Potentially Contributing Sensors" column. At the lowest level of abstraction, observed

TABLE 6.3 Mapping between Sensor Classes and Activities for a Bridging Operation

State	MTI Radar	SAR	COMINT	ELINT	FLIR	Optical	Acoustic
Engineers move to river bank	•		•			•	•
Construction activity		•	•	•	•	•	•
Forces move toward river bank	•		•	•		•	•
Forces move from opposite side of river	•		•			•	•

behavior can be interpreted with respect to a highly local perspective, as indicated in column 6, "Local Interpretation." By assuming that the object is performing some higher level behavior, progressively more global interpretations can be developed as indicated in columns 7 and 8.

Individual battle space objects are typically organized into operational or functional-level units, enabling observed behavior among groups of objects to be analyzed to generate higher level situation awareness products. Table 6.3 categorizes the behavioral fragments of an engineer battalion engaged in a bridge-building operation and identifies sensors that could contribute to the recognition of each fragment.

Situation awareness development involves the recursive refinement of a composite multiple level-of-abstraction scene description. Consequently, the generalized fusion process model shown in Figure 6.3(b) supports the effective combination of (1) domain observables, (2) *a priori* reasoning knowledge, and (3) the multiple level-of-abstraction/multiple-perspective fusion product. The process refinement loop controls both effective information combination and collection management. Each element of the process model is potentially sensitive to implicit (non-sensor-derived) domain knowledge.

6.3 Fusion Process Model Extensions

Recasting the generalized fusion process model within a biologically motivated framework establishes its relationship to the more familiar manual analysis paradigm. With suitable extensions, this biological framework leads to the development of a problem-solving taxonomy that categorizes the spectrum of machine-based approaches to reasoning. Drawing on this taxonomy of problem solving approaches helps to

- Reveal underlying similarities and differences between apparently disparate data analysis paradigms,
- Explore fundamental shortcomings of classes of machine-based reasoning approaches,
- Demonstrate the critical role of a database management system in terms of its support to both algorithm development and algorithm performance,
- Identify opportunities for developing more powerful approaches to machine-based reasoning.

6.3.1 Short-, Medium-, and Long-Term Knowledge

The various knowledge forms involved in the fusion process model can be compared with short-term, medium-term and long-term memory. *Short-term memory* retains highly transient short-term knowledge; *medium-term memory* retains dynamic, but somewhat less transient medium-term knowledge;* and *long-term memory* retains relatively static long-term knowledge. Thus, just as short-, medium-, and long-term memory suggest the durability of the information in biological systems, short-, medium-, and long-term knowledge relate to the durability of the information in machine-based reasoning applications.

---

* In humans, medium-term memory appears to be stored in the hippocampus in a midprocessing state between short-term and long-term memory, helping to explain why, after a trauma, a person often loses all memory from a few minutes to a few days.

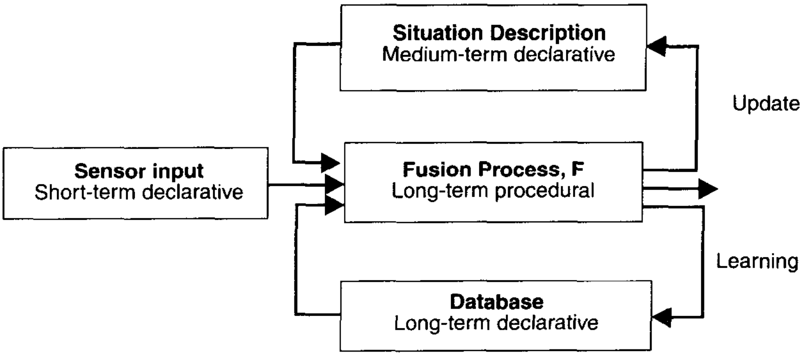

FIGURE 6.7 Biologically motivated metaphor for the data fusion process.

Within this metaphor, sensor data relates to the short-term knowledge, while long-term knowledge relates to relatively static factual and procedural knowledge. Because the goal of both biological and artificial situation awareness systems is the development and maintenance of the *current relevant perception* of the environment, the dynamic *situation description* represents medium-term memory. In both biological and tactical data fusion systems, *current* emphasizes the character of the dynamically changing scene under observation, as well as the potentially time-evolving analysis process that could involve interactions among a network of distributed fusion processes. Memory limitations and the critical role medium-term memory plays in both biological and artificial situation awareness systems enables only **relevant** states to be maintained. Because sensor measurements are inherently information-limited, real-world events are often nondeterministic, and uncertainties often exist in the reasoning process, a disparity between *perception* and *reality* must be expected.

As illustrated in Figure 6.7, sensor observables represent short-term declarative knowledge and the situation description represents medium-term declarative knowledge. Templates, filters, and the like are static declarative knowledge; domain knowledge includes both static (long-term) and dynamic (medium- and short-term) declarative context knowledge; and F represents the fusion process reasoning (long-term procedural) knowledge. Thus, as in biological situation awareness development, machine-based approaches require the interaction among short-, medium-, and long-term declarative knowledge, as well as long-term procedural knowledge. Medium-term knowledge tends to be highly perishable, while long-term declarative and procedural knowledge is both learned and forgotten much more slowly. With the exception of the difference in the time constants, learning of long-term knowledge and update of the situation description are fully analogous operations.

In general, short-, medium-, and long-term knowledge can be either *context-sensitive* or *context-insensitive.* In this chapter, context is treated as a conditional dependency among objects, attributes, or functions (e.g., $f(x_1,x_2|x_3 = a)$). Thus, context represents both explicit and implicit dependencies or conditioning that exist as a result of the state of the current situation representation or constraints imposed by the domain and/or the environment.

Short-term knowledge is dynamic, perishable, and highly context sensitive. Medium-term knowledge is less perishable and is learned and forgotten at a slower rate than short-term knowledge. Medium-term knowledge maintains the context-sensitive situation description at all levels of abstraction. The inherent context-sensitivity of short- and medium-term knowledge indicates that effective interpretation can be achieved only through consideration of the broadest possible context.

Long-term knowledge is relatively nonperishable information that may or may not be context-sensitive. *Context-insensitive* long-term knowledge is either generic knowledge, such as terrain/elevation, soil type, vegetation, waterways, cultural features, system performance characteristics, and coefficients of fixed-parameter signal filters, or context-free knowledge that simply ignores any domain sensitivity. *Context-sensitive* long-term knowledge is specialized knowledge, such as enemy Tables of Equipment, context-conditioned rule sets, doctrinal knowledge, and special-purpose two-dimensional map overlays (e.g., mobility maps or field-of-view maps). The specialization of long-term knowledge can be either fixed (*context-specific*) or conditionally dependent on dynamic or static domain knowledge (*context-general*).

Attempts at overcoming limitations of context-free algorithms often relied on fixed context algorithms that lack both generality and extensibility. The development of algorithms that are implicitly *sensitive* to relevant domain knowledge, on the other hand, tends to produce algorithms that are both more powerful and more extensible. Separate management of these four classes of knowledge potentially enhances database maintainability.

6.3.2 Fusion Classes

The fusion model depicted in Figure 6.3(b) views the process as the composition among (1) short-term declarative, (2) medium-term declarative, (3) long-term declarative, and (4) long-term procedural knowledge. Based on such a characterization, 15 distinct data fusion classes can be defined as illustrated by Table 6.4, representing all combinations of the four classes of knowledge.

Fusion classes provide a simple characterization of fusion algorithms, permitting a number of straightforward observations to be made. For example, only algorithms that employ short-term knowledge are sensitive to a dynamic input space, while only algorithms that employ medium-term knowledge are sensitive to the existing situation awareness product. Only algorithms that depend on long-term declarative knowledge are sensitive to static domain constraints.

While data fusion algorithms can rely on any possible combination of short-term, medium-term, and long-term declarative knowledge, every algorithm employs some form of procedural knowledge. Such knowledge may be either explicit or implicit. *Implicit* procedural knowledge is implied knowledge, while *explicit* procedural knowledge is formally represented knowledge. In general, implicit procedural knowledge tends to be associated with rigid analysis paradigms (i.e., cross correlation of two signals), whereas explicit procedural knowledge supports more flexible and potentially more powerful reasoning forms (e.g., model-based reasoning).

All fusion algorithms rely on some form of procedural knowledge; therefore, the development of a procedural knowledge taxonomy provides a natural basis for distinguishing approaches to machine-based reasoning. For our purposes, *procedural* knowledge will be considered to be long-term *declarative knowledge* and its associated *control knowledge*. Long-term declarative knowledge, in turn, is either *specific* or

TABLE 6.4 Fusion Classes

Fusion Class	Declarative Knowledge Class			Procedural Knowledge
	Short-Term Knowledge	Medium-Term Knowledge	Long-Term Knowledge	
1	•			
2		•		
3			•	
4	•	•		
5		•	•	
6	•		•	
7	•	•	•	
8				•
9	•			•
10		•		•
11			•	•
12	•	•		•
13		•	•	•
14	•		•	•
15	•	•	•	•