Upstream Data Fusion: History, Technical Overview, and Applications to Critical Challenges

Andrew J. Newman and Glenn E. Mitzel

Upstream data fusion (UDF) refers to the processing, exploitation, and fusion of sensor data as closely to the raw sensor data feed as possible. Upstream processing minimizes information loss that can result from data reduction methods that current legacy systems use to process sensor data; in addition, upstream processing enhances the ability to exploit complementary attributes of different data sources. Since the early 2000s, APL has led a team that pioneered development of UDF techniques. The most mature application is the Air Force Dynamic Time Critical Warfighting Capability program, which fuses a variety of sensor inputs to detect, locate, classify, and report on a specific set of high-value, time-sensitive relocatable ground targets in a tactically actionable time frame. During the late 2000s, APL began expanding the application of UDF techniques to new domains such as space, maritime, and irregular warfare, demonstrating significant improvements in detection versus false-alarm performance, tracking and classification accuracy, reporting latency, and production of actionable intelligence from previously unused or corrupted data. This article introduces the concept, principles, and applicability of UDF, providing a historical account of its development, details on the primary technical elements, and an overview of the challenges to which APL is applying this technology.

INTRODUCTION

Sensor and Data Fusion

The term "sensor and data fusion" refers to techniques that combine data from multiple sources to produce new information and inferences and achieve more complete, clear, precise, accurate, and timely estimates of the unknown quantities than could be achieved by the use of a single source alone. Fusion of data from multiple sensors provides several advantages over deriving estimates from a single source (see, e.g., Refs. 1 and 2). First, a statistical advantage is gained by combining independent observations. Second, the relative position or motion of sensors can be exploited to provide a geometric advantage in estimating kinematic states and other attributes of an observed object. Third, the relative strengths and weaknesses of dissimilar data types can be, respectively, magnified and mitigated by combining them judiciously. Fusing different data types from multiple sensor sources, in particular from different sensing phenomenologies (also referred to as modalities), broadens the number of physical observables available to the fusion process, which results in significant improvements in detection, tracking, and classification performance as well as resistance to countermeasures and changing conditions.

The U.S. military has acquired and currently operates a diverse ensemble of intelligence, surveillance, and reconnaissance (ISR) assets but tasks and exploits them in self-contained enterprises, often referred to as "stovepipes," using a combination of automated and manual processes that are highly specialized to a particular ISR asset, data type, application, or military domain and do not systematically interact with other such stovepipes. This approach often fails to fully exploit complementary capabilities and opportunities for collaboration (see, e.g., Refs. 3–6). The sensors provide an enormous volume of data but often do not support precision engagements of targets of interest because of deficiencies in fusion and exploitation of the available data. The process of tasking, collecting, processing, exploiting, and disseminating ISR data is generally divorced from the rapidly evolving tactical picture and the needs of the end user (e.g., theater commander or exploitation system operator). Improved sensor and data fusion capabilities are needed to satisfy the warfighter's accuracy, persistence, and timeliness requirements (see, e.g., Ref. 7).

In general, the ability of a multiple-sensor ensemble to provide persistent coverage on all targets, with high accuracy, high detection probability, and low false-alarm rate, will be constrained by limits on the number and diversity of assets, platform speed, maneuverability, view angle, sensor coverage and update rates, sensor resolution, and mismatch of target and sensor phenomenologies. Coverage gaps, missed detections, false alarms, errors in the sensor support data, and classifier confusion will cause geo-location and classification errors, incorrect associations, track loss or discontinuities, and spurious tracks, resulting in an ambiguous and unreliable tactical ground picture. Fusion and exploitation systems must be robust enough to address these realistic conditions.

Upstream Data Fusion Concept

"Upstream data fusion" (UDF) refers to the processing, exploitation, and fusion of sensor data as closely to the raw sensor data feed as possible within the limits imposed by technical feasibility and operational practicality. Upstream processing minimizes the information loss that can result from the data reduction methods used by current legacy systems that process data from a particular single source; in addition, upstream processing enhances the ability of the fusion process to exploit the complementary attributes of different data sources. The UDF process taps data at an appropriate point in the processing chain near the sensor source (chosen to acquire the desired information content within engineering feasibility and without disrupting current operational processes) and bypasses the data reduction and detection thresholding steps inherent to a traditional single-sensor processing approach. It then exploits the upstream data by tuning detection sensitivity to respond to faint signatures and discriminating between the true targets and the consequent large number of false candidates by fusing data across complementary sensor phenomenologies, diverse view geometries, and different times. In addition, processing of raw (or nearly raw) sensor data allows the UDF algorithms to extract and exploit measurement data (e.g., target position) and associated uncertainties (e.g., error statistics such as variance or covariance) with the highest possible precision, as well as to exploit attribute data that are not normally reported in traditional processing chains. Applying UDF to operational problems in the ground, maritime, and space domains has demonstrated significant improvements in detection versus false-alarm performance, tracking and classification accuracy, reporting latency, and production of actionable intelligence from previously unused or corrupted data. The benefits of UDF are more fully described in the *UDF Benefits* section.

A typical legacy downstream fusion process captures and fuses post-detection data from the available sources. Each of the individual sensor systems applies its own processing, data reduction, and detection thresholding to produce a set of candidate targets (e.g., ground, maritime, or space targets) in its data stream. In particular, each individual system is tuned to optimize its own intrinsic performance, meaning that detection thresholds are set to relatively high levels to maintain a low false-alarm rate while reducing the data processing load. This reduces the probability of detecting targets whose signatures may be only faintly observed in the sensor data (i.e., below the defined threshold). Moreover, data that are identified as "bad" or "corrupted" are usually discarded because they may not be reliable enough to support decision making within that individual processing chain, even though these data can often reinforce (or contradict) the information accrued from other sensors. The performance of a downstream fusion process is therefore inherently limited by the reduced set of thresholded data and the resulting limited number of candidate detections that it receives.

Figure 1 illustrates the concept of tapping and exploiting upstream sensor data to produce information that is not available through current means of processing single-sensor data. The UDF technique captures the raw or partially processed upstream data before decisions are made (i.e., before detection thresholds are applied) and performs an efficient multiple-level screening process to search the upstream data for candidate detections. Screening thresholds are set very low to ensure that data that would otherwise be rejected are considered in the fusion process. This increases the probability of detection of actual targets but also increases the number of false alarms that are passed to the fusion process. The false alarms will typically be rejected in the fusion process, because, in general, there will not be corroborating evidence from other sensor inputs. In this manner, UDF discovers targets and activity that would not be found using a legacy fusion process.

When exploited in their traditional stovepipes, the different data types, or sensing modalities, that are available for a particular mission can produce very different observables, intermediate information products, and end products or releasable report formats. UDF efficiently extracts, accrues, and reports information by using specialized screener, fusion, and output conditioning components (described in the *UDF Design Overview and Primary Technical Elements* section). Figure 2 illustrates an example of applying the UDF methodology to a set of dissimilar sensor modalities (intercepted signals, images, conventional radar, and ground moving target indicator radar) supplying pieces of complementary upstream information that are traditionally exploited separately and combined using only products intended for end-user consumption (if at all).

UDF Principles

UDF is based on the principles of efficient information processing and rigorous model-based evidence accrual for multiple, possibly highly dissimilar, data types. Efficiency is a key tenet because of the need to extract all useful information from a potentially overwhelming volume of source data while maintaining computational tractability at achievable data transmission rates. It is achieved through a modular, distributed architecture in which

- software components referred to as "screeners," specialized to each sensor, process raw data from each sensor individually and transmit a dramatically reduced volume of essential data elements to the fusion process; and
- fusion and output conditioning software components, common to all sensors, exploit upstream information from the screeners to quickly but judiciously eliminate statistically unlikely possibilities.

Mathematically rigorous model-based evidence accrual is a key tenet (*Dynamic Time-Critical Warfighting Capability Feasibility Demonstration Plan*, 23 August 2001) because of the need to properly combine multiple, possibly highly dissimilar, data types in situations where each data type contributes essential information to the fused solution but cannot resolve the targets or their key attributes individually. Each sensing modality produces a distinct set of observables or measured quantities that may not be immediately comparable in form or dimensionality. These must be combined objectively (e.g., without subjective weights) within a common mathematical framework. Inferences must be updated as the data are received and must accurately reflect the accumulated information and aggregate uncertainty.

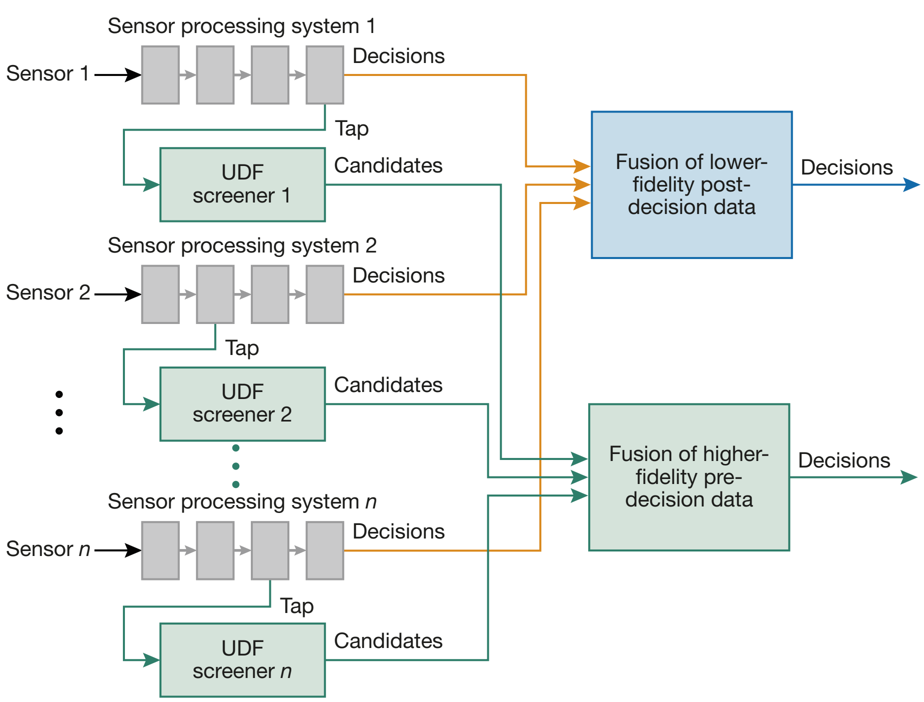

Figure 1. UDF concept for tapping upstream sensor data and bypassing single-sensor processing stovepipes.

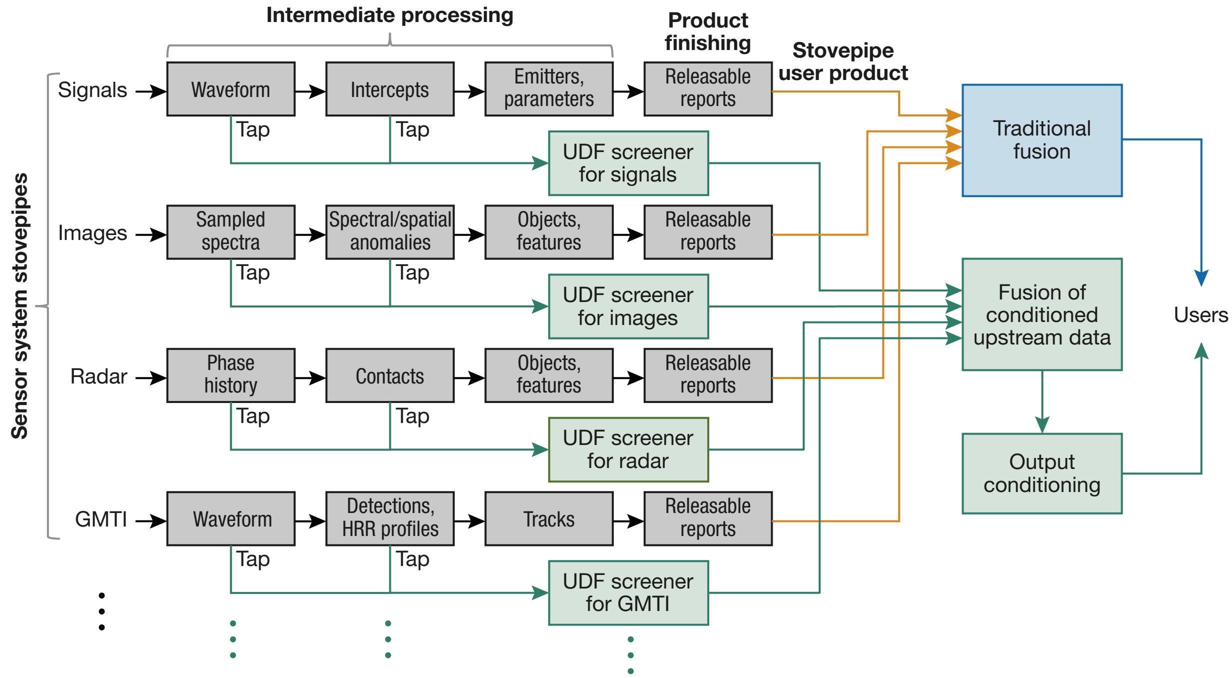

Figure 2. UDF concept applied to several dissimilar sensing modalities. GMTI, Ground moving target indicator; HRR, high range resolution.

The screeners decide what data to transmit to fusion (e.g., candidate detections) and condition those data to prepare them for consumption by a fusion process adhering to the principle of mathematically rigorous model-based evidence accrual (a process referred to as "data conditioning"). The screeners are physically (e.g., geographically) distributed, each as close as possible to its corresponding sensor given practical considerations (e.g., at a ground station or on board the sensing platform). They are also logically distributed in the architecture, meaning that each screener is an independent functional unit that interfaces to the fusion process through a predefined information exchange schema. Each screener is specialized to the sensor type and runs without interfering with the existing operational processing flow. The screeners extract candidate detections with detection thresholds intentionally set low to maximize detection sensitivity so that faint signatures pass reporting thresholds. As a consequence, the screener passes a much larger number of potentially false detections than would normally be tolerated by a traditional detection process reporting targets to a user or downstream fusion process. The screeners transmit only the essential data elements for candidate detections from the upstream processing to the fusion process for correlation and state estimation.

Data conditioning includes some commonly applied sensor data processing functions, such as extracting measurements and features, converting to common units, and aligning to common reference frames, as well as more subtle functions necessary to apply UDF principles. For example, it is critical to the evidence accrual-based fusion process that information supplied by the screener be extracted purely from the source data, unaltered by incorporating exogenous sources or assumptions (referred to as "prior-free" information). (For example, it is traditional and appropriate for legacy sensor systems to incorporate terrain elevation data directly into products intended for end users.) A serious consequence of violating the prior-free prerequisite is the risk of accruing the same exogenous information repeatedly and redundantly in fusion as if it represented new and independent data. Also, for example, screeners typically output engineering support, measurement, and attribute data with fidelity and precision that is not normally passed to an end user but is necessary to associate observations with objects under conditions of high rates of false alarms.

The fusion component (or distributed set of fusion components) performs data association, kinematic state estimation, and class estimation by applying statistical methods to the candidate detections from multiple screeners, over time, to identify candidate detections with true targets of interest, objects that are not targets of interest (e.g., persistent clutter), or random noise in the sensor data. (Current UDF realizations use a centralized fusion component. Decentralized fusion is a subject of current research that is not covered in this article.) The data association process uses a multiple-hypothesis formulation, referred to as Multiple Hypothesis Data Association (MHDA), which applies Bayesian evidence accrual (see, e.g., Ref. 8) to recursively evaluate likelihood ratio statistics on multiple association hypotheses and enable the correct hypothesis to dominate competing incorrect hypotheses over time. The rigorous application of likelihood-based statistical inference requires physics-based and empirically derived models of the sensors (with emphasis on error processes, precisely quantified uncertainties, and sensitivity of measurements to states), targets (with emphasis on statistical distributions of observable features and signatures), and background environment. This approach properly considers a large number of uncertain possibilities without making early decisions or applying heuristics with limited applicability. It exploits the complementary attributes of diverse sensor phenomenologies, sensor geometries, and data collected over time to maintain a low system-level (post-correlation and report release) false-alarm rate.

Output conditioning is a post-processing component that prepares fusion output for use by a human user and determines when the end product is releasable to the user.

UDF History

The development of UDF at APL was initiated by two retired Air Force officers, the late Roy Robinette and the late Michael ("Cisco") Francisco (Fig. 3, left and center, respectively). On the basis of their experience as combat pilots in Vietnam, they were convinced that a great deal of operationally relevant sensor data was not being fully exploited by the military (later documented in Ref. 7). To address the perceived deficiency, they developed a concept they called the "Five-Minute War." The name was chosen to emphasize the potential capability to quickly find and prosecute lethal threats during the short periods of time when they would expose themselves to surveillance. Roy and Cisco presented the idea to industry, but it did not generate significant interest because of skepticism that such a concept could ever be made practical. Consequently, they convened a nonprofit partnership consisting of APL, Georgia Tech Research Institute, and Draper Laboratory to assess the feasibility and utility of the concept. A preliminary conference was held at the Air Force's Unmanned Aerial Vehicles Battle Laboratory, then at Eglin Air Force Base, in January 2000. This was followed by a larger conference at APL in March 2000. At the end of this conference, the three organizations agreed to form a development team with APL as the lead. The concept was discussed and refined at the National Correlation Working Group Workshop, held at Ft. Monroe, Virginia, on 23–25 May 2000.[9]

By August 2001 the development team was convinced that the "Five-Minute War" concept was feasible (documented in *Dynamic Time-Critical Warfighting Capability Feasibility Demonstration Plan*, 23 August 2001). Around this time, the term "upstream data fusion" was coined to succinctly refer to the underlying enabling technology. The team proposed the initiation of a major new research program to develop a prototype capability and test it in a series of progressively more challenging demonstrations. The first step, completed in March 2002, was a more detailed investigation of potential benefits (as described in the *UDF Benefits* section) and demonstration of selected benefits by using previously collected sensor data.

Following that success and under the leadership of a visionary APL program manager, Bill Walker (Fig. 3, right), a multidisciplinary team was formed by adding 12 new industry partners to the original three, bringing the total to 15 organizations and more than 200 individuals. The new partners were BAE Systems, Defense Consultants Ltd., Dynetics, Francisco & Associates Inc., Keith S. Peyton, Lockheed Martin Corporation, Orincon, Science Applications International Corporation (SAIC), Scitor Corporation, Titan Corporation – Aerospace, Veridian Systems Division Inc., and Zel Technologies, LLC. The team set a goal of conducting a live end-to-end demonstration by May 2004. To handle the complexity and risk associated with the proposed effort, the team made two key decisions. First was the adoption of an agile software development method. Because the concept was so new, no firm requirements existed. Traditional software development presumes that the technology and requirements are available at the outset and are stable throughout development. Agile methods accommodate frequent adaptation and were, therefore, well suited to development of the UDF prototype. After careful comparison of the available options, an agile software development method known as "Scrum" (see, e.g., Ref. 10) was selected. Second, the team adopted a philosophy that emphasized cooperation over competition, wherein the prototype was synthesized by choosing, modifying, and integrating pieces of algorithms and code contributed by team members. The team referred to the cooperative synthesis approach as "breed-the-best," which is a

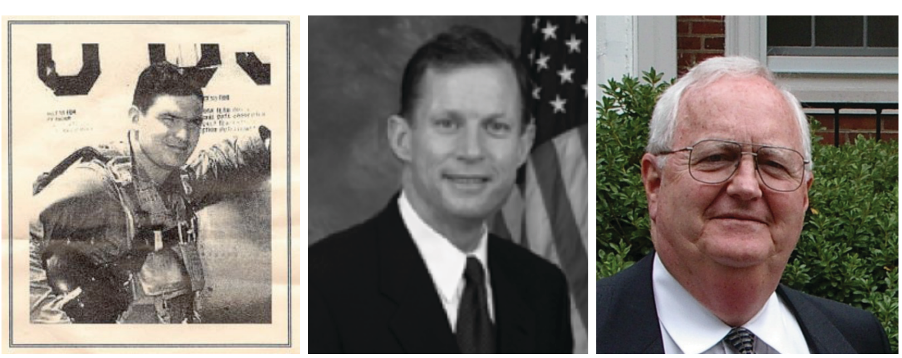

Figure 3. UDF pioneers Roy Robinette, Michael ("Cisco") Francisco, and Bill Walker. (Michael Francisco photo courtesy of *Air and Space Power Journal*.)

word play contradiction on the commonly used "best-of-breed" approach whereby synthesis is achieved through competition by choosing the best finished product or component for each individual function. The development philosophy proved surprisingly effective in producing a working prototype within the schedule and funding constraints, leading the team to affectionately call the resulting product "stone soup."

This was truly a team effort, as shown by the many contributors across a number of organizations (see *Acknowledgments*). A few members of that team and its successors are mentioned in what follows with the sincere hope that neither the authors nor the readers will diminish the contributions of others. (All contributors are or were affiliated with APL unless otherwise noted.) Dr. David Porter, Dr. James Christ, and Dr. Glenn Mitzel developed the original functional architecture design (an early, more detailed version of Fig. 1) in an intense 2 weeks of isolation. It was subsequently refined by the larger team and became the framework for all other development. The physical architecture was designed under the leadership of Glen Long. The software architecture was developed under the leadership of Steve Wieprecht. The team applied the Scrum software development process under the leadership of Will Menner. The screeners for two of the sensing modalities were developed under the leadership of Dr. Christopher Boswell, Dr. Patricia Murphy, and Dr. Chung Fu Chang (Lockheed Martin). One screener design was based on prior work by Myron Brown and extended to include new target features and detection and geo-location techniques. The target features and the complex code required to extract them were contributed by multiple organizations, each having decades of experience with the specific sensing modalities. For another sensing modality, a relatively mature automated fusion capability was available and adequate for the first live demonstration. Applying the breed-the-best philosophy, that capability was modified to accept raw data from other screeners under the leadership of Marshall Alworth and Mike Rector (both with Scitor Corporation; Alworth has since joined APL). Output conditioning, which turned the impossibly complex output of the automated fusion capability into a final end-user product, was developed by Jeff Gilbert.

The demonstration in May 2004 was highly successful. Real targets were deployed on test ranges where their positions were precisely measured for comparison with reported locations. During the final preparations for the demonstration, after one long day of work and with initial installation at the demonstration site complete, the site team turned on the prototype and was ready to leave for the night. Suddenly the first live data appeared and the prototype processed the data. The site team stayed for hours convincing themselves that the results seemed reasonable. That event became known to the whole team as the day "Skynet went operational" (Skynet refers to a fictional revolutionary artificial intelligence system that gains self-awareness in the 1984 movie *The Terminator*). Later, on the day of the live demonstration, raw data collected by operational sensors were combined automatically by the prototype. At one instant, with local and remote visitors watching the unscripted events, the combined results showed a dramatic improvement over observations seen by one sensor alone. Remote visitors were not able to see the results immediately at that moment, but they did note the sudden eruption of clapping and shouting that they heard coming from team members at the demonstration site.

After the successful May 2004 demonstration, another demonstration was planned for June 2006 under more stressful conditions, including some target countermeasures. For this demonstration and to support future development, an entirely new fusion algorithm was designed by Dr. David Porter and the late Dr. Larry Levy. This algorithm was designed from scratch to employ fully recursive logic and to treat ingested modalities in a balanced and mathematically rigorous fashion. The algorithms were implemented in software by John Florence and Dr. Andrew Newman. Additional improvements were made to the data screeners and to the overall software architecture. The demonstration in June 2006 was, again, highly successful. Despite the countermeasures, targets were reliably detected and accurately located with very few false alarms. At one point, the Air Force Chief of Staff was invited to watch the live results at the Pentagon. Under the unscripted conditions, very little was happening during the short period he was able to visit. However, moments after he left a very successful event occurred. His subordinates quickly sent him an e-mail message reporting the results. Soon afterward UDF became known supportively as "dumpster diving" and "Trash-INT" to emphasize how it could glean useful information from collected data that would otherwise be discarded.

The next live demonstration occurred in June 2008, which successfully showed the capability to process additional sensors and report on additional target types. In March 2009 the prototype was activated for the first time with continuous operational feeds. This test was conducted over a 2-week period, primarily to assess robustness. Operational availability, i.e., the fraction of time the prototype ran unattended, was extremely high. The prototype was deployed as a limited operational capability for nearly 3 months in the fall of 2009 and on several occasions since. During the test and assessment process, the prototype has been undergoing continuous refinement and maturation based on these operational experiences. An initial prototype deployment for use by intelligence analysts at operational sites is planned for the spring of 2013.

Since 2008, APL has expanded the application of UDF techniques to critical challenges and new domains including space, maritime, and irregular warfare, described in the *UDF Application Areas* section. In addition, APL has recently initiated internal research and development efforts to investigate the feasibility of applying UDF technology to detection and characterization of underground facilities, protection of docked submarines carrying nuclear weapons, and area defense against terrorist attacks.

UDF BENEFITS

Detection

The improved detection capabilities of UDF can be illustrated using a simplified analysis (originally documented in an APL internal presentation by Mitzel entitled "Illustration of the Potential Benefits of Upstream Data Fusion for Target Detection," dated April 2001) with two notional sensor systems with identical processing, decision logic, and performance characteristics. For the purposes of this analysis, the UDF concept simply depicted in Fig. 1 is further simplified to two sensors where the sensor processing for each sensor is assumed to consist of computation of a decision variable and the application of decision logic to that variable. Figure 4 depicts three distinct approaches for determining whether a target is present or absent given the source data from two sensors. First, decisions are produced from each individual sensor using its intrinsic processing and decision logic (orange). Second, one step upstream (blue), a decision is produced by combining the decisions made by the two individual sensor processing chains. This approach may require some modification of the individual sensor decision logic to adjust the individual decision thresholds for the purpose of maintaining a fixed probability of detection. Third, two steps upstream (green), a decision is produced using the internal decision variable of each of the two individual sensor processing chains.

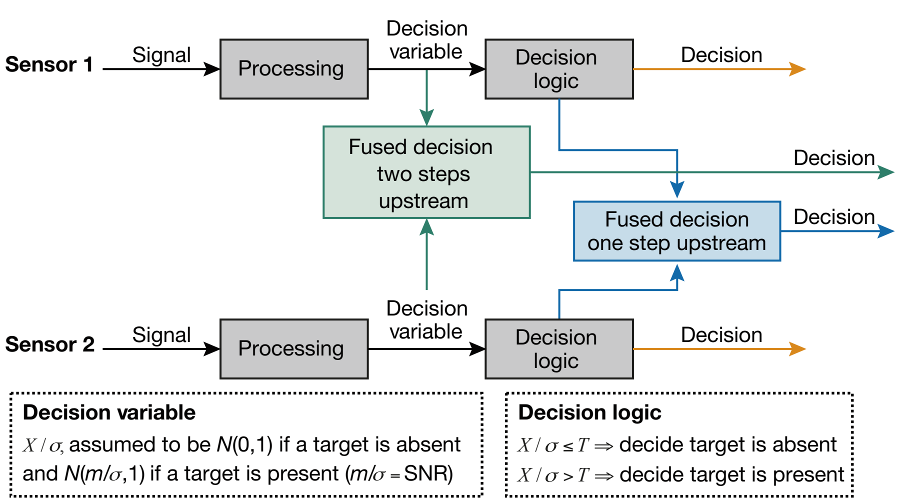

Figure 4. Fusion decision logic applied at varying steps upstream.

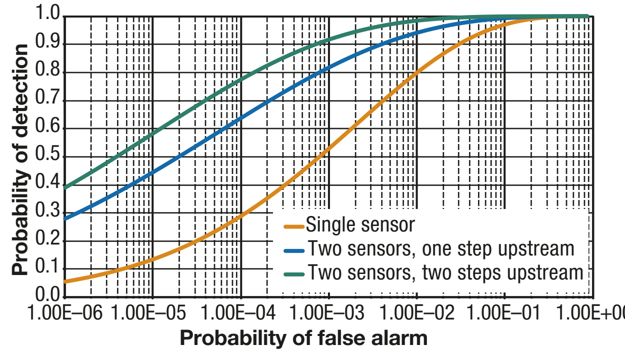

Figure 5. Comparison of ROC curves for single sensor and UDF one and two steps upstream. SNR = 10 dB.

It is common to characterize sensor system performance by using a receiver operating characteristic (ROC) curve, which plots probability of detection versus false alarm rate over a range of detection thresholds. Figure 5 shows the performance gain when going from single-sensor detection processing (orange) to one-step upstream fusion (blue) in which a Boolean AND operation is applied to combine the two single sensor decision logics, and further to two-steps upstream fusion (green) in which the decision variables are accessed and combined apart from individual decision logics by using a joint probability distribution function.

The single-sensor ROC comes from classical binary detection theory. For normally distributed decision variables it can be expressed as

$$P_{d,k} = 1 - \Phi\left[\Phi^{-1}(1 - P_{f,k}) - {}^{m_k}\!/_{\sigma_k}\right],$$

where k = 1,2 for the individual sensors; Φ is the cumulative distribution function of a zero-mean, unit-variance normal random variable; m is the decision variable mean (imposed only by the target); σ is the decision variable standard deviation; P_d denotes probability of correct detection; and P_f denotes probability of false alarm. The signal-to-noise ratio (SNR) is therefore given by m/σ. The orange line in Fig. 5 plots this curve for a common SNR of 10 dB (roughly 3.16).

For this analysis, the logic rule for fusing decisions one step upstream is defined to require that the sensors agree on detection decisions. Other logic rules (such as those requiring at least one detection) are possible but are not considered here. To achieve the same probability of detection for the fused result as for the individual sensors, the individual sensor probabilities of detection must be increased by lowering the thresholds. Assuming that the individual sensor statistics are independent, the composite probability of detection is the product of the individual probabilities of detection. Assuming that the decision variables for both sensors are identically distributed, the ROC one step upstream is given by

$$P_{d,1\text{ step upstream}} = \{P_d\}^2 = \left\{1 - \Phi\left[\Phi^{-1}\left(1 - P_{f,1\text{ step upstream}}\right) - \frac{m}{\sigma}\right]\right\}^2 .$$

The blue line in Fig. 5 plots this curve for a common SNR of 10 dB.

For fused decisions two steps upstream, this analysis assumes simultaneous stimuli of the two sensors but with access to the individual decision variables themselves. The new decision logic requires the sum of the two normalized decision variables to exceed a threshold. It is easy to see that such decision logic may show improvement over one-step upstream logic. For one-step upstream logic, both sensors must meet individual detection thresholds. For two-step upstream logic, the individual thresholds are irrelevant. A weak decision variable measurement in one sensor can be compensated by a strong measurement in the other variable as long as the sum meets the specified sum threshold. The sum is normally distributed with zero mean and a standard deviation of $\sqrt{2}\ m/\sigma$. Therefore, using the same reasoning as in the single-sensor case but with the summed decision variable,

$$P_{d,2\text{ steps upstream}} = 1 - \Phi\left[\Phi^{-1}\left(1 - P_{f,2\text{ steps upstream}}\right) - \sqrt{2}\,\frac{m}{\sigma}\right].$$

When comparing this with the ROC for a single sensor, it is apparent that this is equivalent to a gain in SNR of $\sqrt{2}$ (~3 dB). The green line in Fig. 5 plots this curve for a common sensor SNR of 10 dB.

For a fixed probability of detection it is obvious that the false alarm rates will be substantially reduced as the fusion decision logic is applied further upstream. For instance, at a detection probability of 0.80 the false alarms are reduced by more than a factor of 10 at one step upstream and nearly 100 at two steps upstream. Conversely, for a fixed false alarm rate, the detection probability will be improved. This implies that the system will achieve detections that would otherwise go unreported. Although the illustration is simplistic, detection gains of these types have been repeatedly demonstrated in the field.

Other Benefits

Target Location and Classification

Fusing location estimates from multiple independent sources generally improves accuracy. But the combination of data from multiple sources inevitably requires assumptions about such matters as the independence of errors, temporal error behavior (e.g., from persistent biases), and error magnitudes. When the combination is purely downstream, those assumptions are often made without understanding or knowledge of the underlying error mechanisms.

UDF facilitates access to higher-quality engineering support data. It therefore enables better error accounting and error modeling. Specific examples of the improvement include the elimination of conservative "fudge" factors (ad hoc adjustable parameters with weak theoretical justification), isolation of common error contributions (e.g., elevation data), incorporation of correlation durations for slowly varying biases (systematic errors in the sensor measurements), and the ability to estimate and compensate for bias errors by correlating across different sensing modalities without relying on off-line calibration. Similarly, UDF facilitates access to raw data where subtle features for target classification can only be exploited.

Timeliness

UDF bypasses some of the downstream processing for individual sensors (as shown previously in Fig. 2), which typically includes processing that primarily conditions the individual sensor output product for human consumption. This reduces system processing latency. Furthermore, the UDF process is implemented recursively, so that a very small amount of additional data and processing from one sensor may become the tipping point for reporting given all the other previous data.

Improvements in timeliness can be dramatic. For example, for a single sensor, the number of observations to reach a reporting threshold may be large and may take a significant amount of time to collect. However, in the UDF approach, if a target has been barely missed in previous sensor observations but the memory of that detection has been recursively retained (consistent with the principle of evidence accrual described in the *UDF Principles* section), then the first indications of the target from another sensor, although themselves not reportable, may result in a combined detection long before enough observations from the single sensor would have been collected and exploited.

Robustness to Countermeasures

UDF is a logical response to countermeasures that are designed to reduce target observables, to confound a limited set of modalities, or to restrict the exposure of key target signatures. UDF exploits more subtle target features and scales to include more modalities. Moreover, as explained in the previous section on timeliness benefits, even the most fleeting target signature may provide a sufficient amount of corroborating evidence to detect and report the target.

Required Computational and Communications Capacity

By deploying screeners as far upstream as possible and judiciously controlling the release of fused information to the user according to rigorously computed uncertainties and confidence factors, the UDF framework allows a significant reduction in the communications capacity required to transmit data among networked components and to end users. This reduction is particularly dramatic for sensing modalities producing extremely high volumes of data such as full-motion video (FMV) and wide-area imaging. End-to-end reductions in required communications capacity of many orders of magnitude have been demonstrated in field tests. In addition, the overall computational load is distributed and the processing capacity of each individual component can be chosen according to its specific loading profile.

Asset Employment Efficiency

UDF not only improves awareness of those targets for which it is certain, but it also keeps careful track of the ambiguities and uncertainties about targets that are not yet reportable. This additional information facilitates the optimal use of limited surveillance assets to collect the last piece of accruable evidence as opposed to initiating a new search. Also, accurate assessments of uncertainty and ambiguity form a basis for choosing the assets and focusing their attention where it will be most useful. This UDF benefit forms the technical foundation for a more comprehensive concept referred to as "Closed-Loop Collaborative ISR" (CLCISR), which is the subject of the article by Newman and DeSena in this issue.

UDF DESIGN OVERVIEW AND PRIMARY TECHNICAL ELEMENTS

UDF Processing General Framework

The general framework for UDF processing consists of several distributed upstream screeners performing detection and data conditioning functions; one or more fusion elements performing data association, kinematic state estimation, and class estimation functions; and an output conditioning element, each of which is described in general terms in the sections below (some references are provided for detailed descriptions, although in many cases the details are program specific and access may be limited). These components, and to a lesser extent the framework itself, must be specialized for the particular sensor types and mission while adhering to the general UDF design principles. Current realizations employ a centralized fusion element, although decentralized fusion designs are feasible (see, e.g., Ref. 11) and are the subject of current research. The general framework is illustrated in Fig. 6, which shows one possible realization for a ground target surveillance and reconnaissance mission.

Screeners

Each screener ingests data directly from the source, or with the minimum intervening transmission and processing that is feasible, practical, and advantageous to system performance. The screener often resides on the sensor platform or the processing element that it directly feeds in its traditional processing chain. Screener processing is specialized for the particular data type, using detailed physics-based and empirical models of the specific sensor phenomenology. Therefore, screening algorithms must be developed for each distinct sensor type, and a screener software component is required to process the upstream data from each individual sensor.

The essential functions of each screener are to extract candidate targets from upstream sensor data, measure observables, convert all observable data to a common format and units, align data to common reference frames where necessary, estimate measurement error with the maximum possible precision, and assure that information produced is prior-free. High precision prior-free observable data are essential for fusion to accurately and faithfully compute and represent estimate uncertainties and association probabilities.

In general, each screener may also perform pattern classification algorithms using the extracted feature data to produce conditional likelihoods of the target belonging to particular object classes (e.g., target types, persistent clutter, and random clutter). However, the screener does not perform a traditional target recognition or classification function. It produces only candidates and conditional likelihoods rather than making final decisions about target class or identity. Feature distributions for objects of interest can be estimated as Gaussian mixture densities by applying maximum likelihood expectation maximization and data mining techniques to empirical data. In addition, the screener may also provide the functionality to deal with corrupted measurements and dynamically calibrate sensors to remove bias error.

To keep pace with potentially high volumes of sensor data, low-complexity design of screener algorithms and models for fast computational processing is essential. The techniques and algorithms must be robust to varying conditions (e.g., illumination, collection geometries), be adaptive for new targets and backgrounds, and exploit all available information (e.g., spatial, temporal, and spectral). Screener design often applies a multilevel screening algorithm to efficiently search the upstream source data in near real time to find candidate targets in the data stream. For example, the first-level screening may apply a fast, coarse anomaly-detection technique to identify a very large number of candidate targets. Second and higher levels of screening, as appropriate, progressively screen the remaining candidates, applying more sophisticated algorithms such as pixel-level, object-level, and feature-level algorithms (see, e.g., Ref. 12).

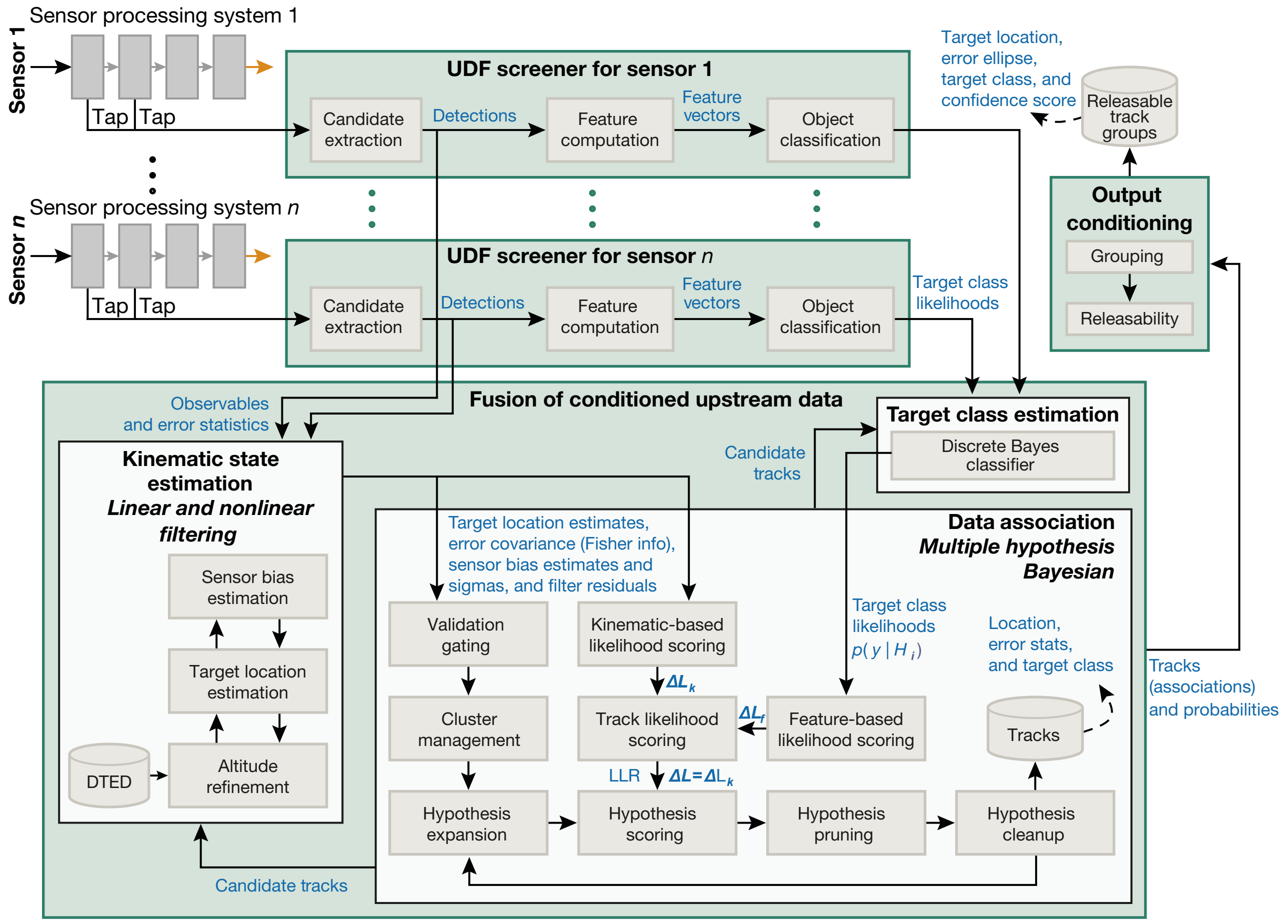

Figure 6. One possible realization of a general framework for UDF processing. DTED, Digital terrain elevation data; LLR, log likelihood ratio.

Most importantly for UDF, the screening thresholds applied by screening are intentionally kept low to maintain a high probability of detecting actual targets of interest. As a consequence, the screener passes a much higher number of detections, which initially produces a much higher false alarm rate than would be tolerated by a traditional detection process reporting targets to a user or downstream fusion process. The UDF framework relies on data association across sensor phenomenologies, sensor geometries, and data collected over time to maintain a low system-level (post-correlation and report release) false alarm rate.

Fusion

The fusion element (or elements) performs data association, kinematic state estimation, and object class estimation. Data association refers to the process of determining the origin of each detection produced by the screeners as an object already observed (and tracked), a new object not yet observed, or a false alarm induced by random phenomena. This includes correlation across the various sensor types and over time. The result is a complete partition of the set of screener detections into subsets corresponding to real objects in the observed physical space or spurious noise (false alarms). Each possible such partition is referred to as an association hypothesis (see discussion of multiple hypothesis techniques below). Kinematic state estimation typically refers to the process of inferring the position and motion variables (e.g., geo-location, velocity, or orbital elements) of detected objects. Object class estimation refers to the process of inferring the types of detected objects chosen from a set of discrete categories. In certain applications, at a higher fidelity, target identity may also be inferred. These functions within the fusion element are tightly coupled, as described below, and are illustrated by example in Fig. 6.

The data association and state estimation functions process candidate detection data from all sensor screeners asynchronously. This is possible because the screeners are designed to report observable data by using common data formats and units known to the data association and state estimation functions. However, the techniques applied depend on expected target properties and behavior.

The multiple hypothesis methodology is essential to data association in the UDF framework where screener detection sensitivity, and consequent false alarm rates, is set intentionally high. The MHDA approach [often referred to as Multiple Hypothesis Tracking (MHT) in the literature because of its common application to moving target tracking problems[13–16]] applies Bayesian evidence accrual to recursively evaluate likelihood ratio statistics on multiple association hypotheses and enable the correct hypothesis to dominate competing hypotheses over time. As the term suggests, "evidence accrual" refers to a technique for combining information from different sources and over time in which estimates and decisions are updated according to physical and statistical models whenever new data become available. Bayesian evidence accrual, or Bayesian inference, applies Bayes's law relating prior and posterior probability densities of the estimated parameters. It provides a straightforward and theoretically sound method of recursively updating belief about unknown random variables by incorporating information from different sources. A likelihood ratio test is a common statistical test used to compare the fit of two models to the available data. The log likelihood ratio is the test statistic that quantifies the relative fits. For MHDA, the log likelihood ratio statistic is computed for each association hypothesis under consideration, with the hypothesis that all detections are false alarms serving as the basis of comparison (see Refs. 14 and 15 and references therein for more details).

The MHDA methodology generates alternative data association hypotheses whenever the data cause an ambiguous or conflicting situation (illustrated by example in Fig. 7). Each hypothesis is a complete and nonconflicting set of detection-to-track associations. Decision making is delayed by propagating hypotheses in anticipation that subsequent data will resolve the uncertainty. Association hypotheses are evaluated, ranked, confirmed, and denied via recursive computation of likelihood statistics derived from the input data and model parameters. As each new set of measurements is made available to the data association process, the algorithm considers the possible ways that the new measurements can associate with the existing tracks. Using precise kinematic information, and based on the consistency of the measurements with the target (or background) models, the system uses Bayesian likelihood techniques to calculate likelihood measures for the hypotheses. Unlikely hypotheses are pruned, and the probability mass is re-normalized over the remaining hypotheses. In this manner, true targets are reinforced by evidence from multiple sensor modalities and will pass the reporting thresholds; false targets decorrelate and are rejected.

The data association function operates in conjunction with the state and class estimation functions. The state estimation function estimates the kinematic state of all candidate target objects. It may also be used to dynamically estimate observable sensor bias errors where applicable. Recursive linear and nonlinear filtering techniques are used to update target and bias state estimates with new measurement data and, depending on the application, to propagate estimates forward in time to correlate with the next set of available measurement data. In general, the MHDA algorithms must properly account for cross-track correlations in the bias states. The class estimation function applies Bayes's law and the law of total probability to update a discrete probability distribution on the class of each candidate target object (see, e.g., Refs. 17 and 18). The MHDA computes the log likelihood score for each candidate association of observations (that form a track) recursively with a data update that combines a kinematic component derived from the filter residuals and residual covariance (or square root information quantities) and a feature-based component derived from classifier-generated conditional probabilities of detected objects in feature space.

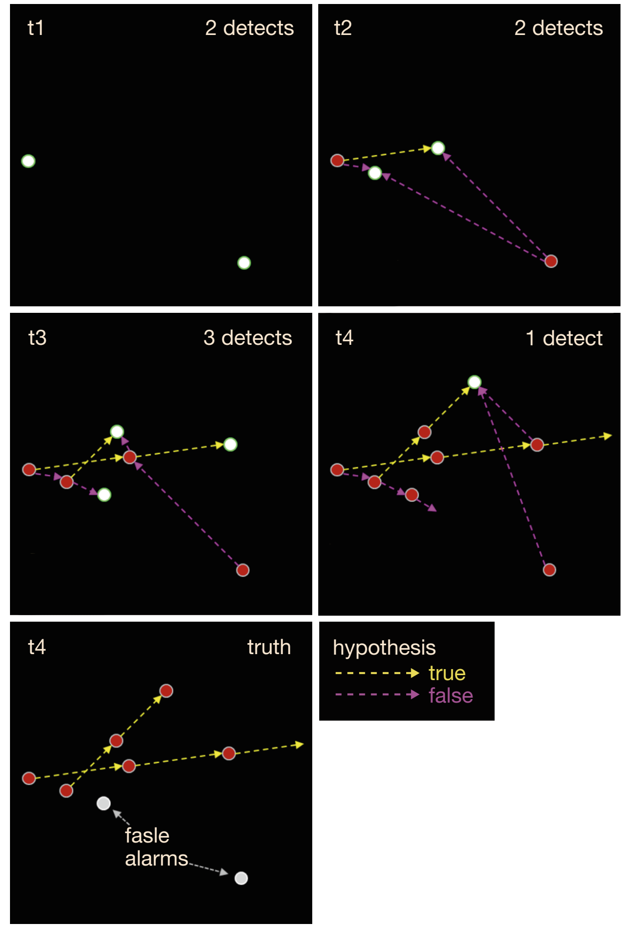

Figure 7. Example MHDA association of detections into tracks over four time steps.

The UDF prototypes developed by APL have been based on the Reid MHT formulation[13] and modern variants.[14–16,19] Traditionally, MHT couples MHDA with independent track filters, commonly Kalman or iterated extended Kalman filter variants that are efficient recursive estimators of the state of a dynamical system from a series of noisy measurements. The APL prototypes instead use an iterated extended square root information filtering (SRIF) technique[19,20] for the track filters to recursively estimate target and bias states from sensor measurements and altitude data.[21] The SRIF algorithm has advantages over the Kalman filter and its other variants in terms of numerical stability and precision, computational efficiency, and amenability to distributed processing. The MHDA and SRIF implementations estimate and compensate for sensor bias errors within the core algorithms, including approximate accounting for cross-track correlation (internal presentations and private communications, D. W. Porter and L. J. Levy, APL, 2003–2006).

The APL prototypes have been designed to run in a multiple-sensor, multiple-target setting where the sensor data may include a mixture of dissimilar data types. The fusion processing is augmented with support for processing ambiguous, nonlinear, and biased measurement,[22] as well as processing of dissimilar data types with performance that is robust to different orders of the input data.[23] The kinematic state estimation function employs and switches between linear and nonlinear filtering techniques at different steps in the algorithm as needed to balance the need for estimation accuracy with the need for computational tractability. The various realizations that have been developed at APL are equipped with specialized capabilities to process and fuse specific data types such as angle measurements from screening optical imagery, range and range rate measurements from radar sources, and other data types relevant to ground and maritime surveillance applications such as time difference of arrival from RF signal intercepts and moving target indicator (MTI) contacts from radar and video sources. Some realizations also include a multiple model capability for tracking targets exhibiting different types of motion behavior.

Output Conditioning

Output conditioning refers to the process of aggregating and resolving competing hypotheses to provide target reports to the end user. Fusion and output conditioning are tightly coupled and require a coherent design to function properly together. The MHDA approach inherently presents situations in which two or more competing hypotheses are similar, cannot be fully resolved given the available data, and the remaining differences are not important for the given application. Presenting these close, but distinct, hypotheses to a user would be confusing and distracting and would reduce the operator's confidence in system performance. The output conditioning function (internal presentations and private communications, J. M. Gilbert, APL, 2003–2011) resolves this ambiguity in the MHDA output by statistically aggregating similar association hypotheses to produce logical entities, called "groups," that represent possible targets. These groups are the main objects of interest in determining releasability of target reports to the tactical user and computing metrics on fused output. Output conditioning provides a rigorous means of aggregating two or more hypotheses or target tracks that individually are too ambiguous or have insufficient confidence to pass a releasability threshold but can do so in aggregate.

UDF APPLICATION AREAS

Relocatable Time-Sensitive Ground Targets

The most mature application of UDF is the Air Force Dynamic Time Critical Warfighting Capability (DTCWC) program currently managed by the Air Force Research Laboratory (AFRL). DTCWC applies UDF to a variety of sensor inputs to detect, locate, classify, and report on a specific set of high-value, time-sensitive relocatable ground targets in a tactically actionable time-

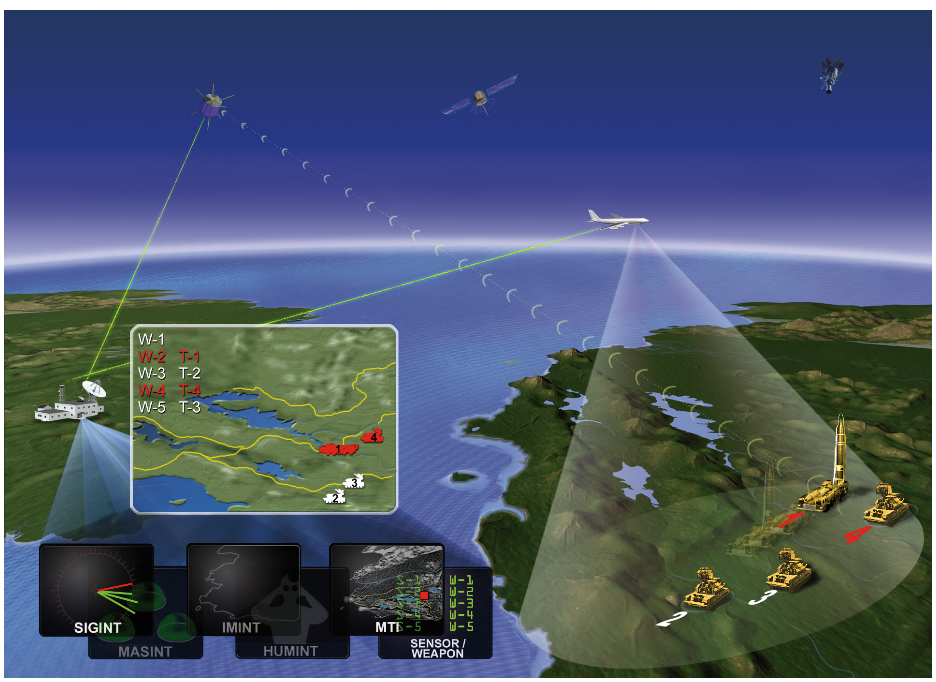

Figure 8. UDF of diverse sensors against relocatable ground threats.

frame.[24–27] Figure 8 illustrates a prototypical mission supported by DTCWC. An automated UDF capability is the critical enabling technology for DTCWC because of the need to prosecute evasive targets employing countermeasures (e.g., frequent relocation, camouflage, and other mitigation of observables and spoofing of signatures) during short periods of time when they expose themselves to surveillance. DTCWC has successfully demonstrated an automated end-to-end capability to fuse live data feeds in several controlled field demonstrations on test ranges (May 2004, June 2006, and June 2008) as well as in uncontrolled limited operational capability (LOC) assessments using operational sensor data against real targets in operational environments (several during 2009–2012).

The early history of what was to become DTCWC, including some details about the early field demonstrations, is described in the *UDF History* section. The 2004 field demonstration validated the technical feasibility of the UDF concept as applied to the DTCWC mission. The 2006 field demonstration showed the feasibility of the MHDA-based fusion approach, the capability to perform successfully in a more challenging natural environment under more realistic conditions including some target countermeasures, and the capability to provide timely actionable target information to users in a targeting cell.

In June 2008, DTCWC was exercised over multiple days by using live national and tactical sensor data feeds to demonstrate and assess new capabilities including integration of a new tactical sensor, reporting on a new target of interest, a reengineered fusion engine, and a browser-based user interface. DTCWC was operated from an Air Force distributed ground system site, which controlled the distributed system components resident at data collection sites. DTCWC achieved reporting accuracy, a false alarm rate, an estimate error, and performance latency that were all within operational parameters for tasking the new tactical sensor in its tipping and cueing role.

In March 2009, DTCWC was run in an operational setting for the first time, with the primary objective of assessing system availability for users. The prototype ran nearly continuously for more than 2 weeks. With that success to bolster the case, in the fall of 2009 DTCWC was assessed at a Technology Readiness Level of 7 by an Air Force-led, independent assessment team. The results of these UDF demonstrations and assessments were recognized and documented by the Joint Defense Science Board Intelligence Science Board Task Force on Integrating Sensor-Collected Intelligence.[28]

From October 2009 through January 2012, the DTCWC program conducted a series of LOC assessments during which DTCWC processed live operational data from supported sources and produced detection reports within specified areas of interest. DTCWC detection reports were made accessible to trained personnel for use in supporting operations, collections, and analytic decisions. During 2012 and 2013, DTCWC has been deploying an initial prototype for use by intelligence analysts at operational sites as the first step in transitioning to an operational capability.

The field demonstrations and LOC assessments have confirmed the feasibility and operational utility of UDF in support of the DTCWC mission against a set of relocatable time-sensitive ground targets. They have shown that UDF can achieve high detection probability (e.g., detecting targets that would have been missed by a single-sensor system) with very low false alarm rates, and geo-location accuracy capable of supporting target engagement. Furthermore, they have shown the ability to overcome corrupted data, whereby information from one sensor system was used to automatically reject poor measurements from another system, and to automatically estimate and remove certain types of bias errors.

Maritime Situational Awareness

The Navy needs to track ships globally to counter resurgent and emerging surface naval threats.[29,30] Worldwide ship tracking typically relies on a combination of marine radar and vessel self-reporting, such as the Automatic Identification System (AIS), a system whereby ships broadcast their unique identification, position, course, and speed. However, these sources of information are not adequate for tracking many important vessels that are of interest to the Navy and may be attempting to hide in regions of dense maritime traffic, conceal their identity and position, or otherwise defeat surveillance.

The maritime surveillance problem presents unique challenges in terms of the extremely large area required for coverage, long track durations, and high density of vessels in many areas of interest (Fig. 9). Available sensor types include surface MTI radars, a variety of imaging modalities, and passive systems such as AIS and RF signal receivers. However, there is no single sensor that detects, tracks, and identifies ships throughout their voyages over wide areas.[31] Moreover, individual sensors suffer from weaknesses such as periodic coverage, narrow coverage, unreliable detection, inaccurate kinematic information, imprecise identification, and high sensitivity to weather conditions. It is necessary to combine and correlate data from disparate surveillance sensors. This is further challenged by dense shipping backgrounds and unpredictable latencies in data transmission that cause data from different sources to arrive out of sequence relative to collection time.

In a series of efforts using real and simulated sensor data, APL showed the potential of UDF to improve maritime situational awareness by screening and fusing upstream data from a variety of operational sensors that

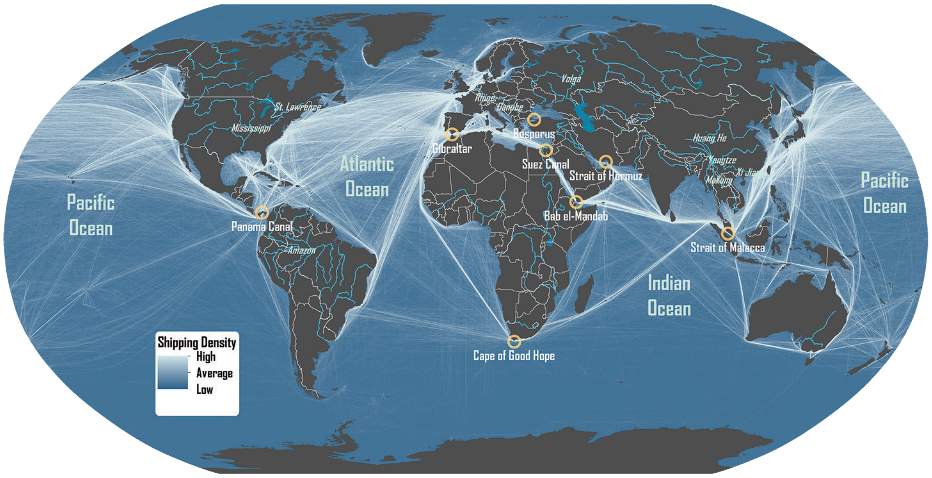

Figure 9. Global maritime ship traffic density. (Reproduced with permission from "Maritime Shipping Routes and Strategic Passages," *The Geography of Transport Systems*, http://people.hofstra.edu/geotrans/eng/ch1en/appl1en/maritimeroutes.html, © 1998–2012, Dr. Jean-Paul Rodrigue, Hofstra University, New York.)

are currently exploited separately, and to do so without interfering with their current operational procedures or capabilities. These efforts emphasized the benefit of fusing sensors that provide strong target feature information with sensors that provide persistent area coverage and precise target location information. They also showed the adaptability of the UDF framework, with relatively minor modifications to current realizations, to maintaining track custody of maneuvering surface targets over long periods of time while traveling through dense backgrounds.

In this work, track-oriented MHT techniques[16] were applied to more efficiently process sensor data for long-duration tracking (by contrast, the DTCWC time-sensitive targeting mission forced an emphasis on track initiation for targets with short exposure times). Multiple motion model filtering techniques (see, e.g., Ref. 32) were applied to enhance tracking of maneuvering ships. A novel nonlinear filtering technique, using an iterated extended SRIF, multiple motion models, and the Dyer–McReynolds smoother,[33] was developed to enable tracking of maneuvering ships using out-of-sequence measurement data. The results showed that, even in dense ship traffic environments, the UDF approach can reinforce the existing sensor track picture.[34] This improvement could provide valuable intelligence to the warfighter on existing tracks and enable cues to vessels of interest entering the platform's area of regard.

Space Situational Awareness

As our reliance on space systems and space technology grows, so does the potential for collisions and other spacecraft interactions, leading to potentially serious consequences. It is essential to maintain an accurate catalog of most, if not all, objects in orbit (Fig. 10). Traditional space surveillance and tracking systems operate with high detection thresholds with an emphasis on rejecting false alarms. The post-detection data are fused to produce orbit tracks. In addition, processing and exploitation of metric data (e.g., satellite position and velocity) and feature data (e.g., photometric and radiometric signatures) are traditionally performed in separate organizational stovepipes. Traditional methods have performed well in finding, tracking, and identifying larger objects in Earth orbit, but smaller objects, including a variety of potential threats, are often lost or missed entirely. In addition, current methods are severely challenged when satellites maneuver from their expected orbital trajectory or are in close proximity to (and hence possibly confused with) other satellites. Enhancements to persistence and timeliness of space situational awareness (SSA) are needed to deal with stressing cases such as dim, closely spaced, and maneuvering objects.

In 2009–2010, APL completed a project sponsored by the Air Force Space and Missile Systems Center to investigate the operational utility of UDF techniques to detection and tracking of dim resident space objects (RSOs) and discrimination of closely spaced RSOs using space surveillance sensor data. The project executed a

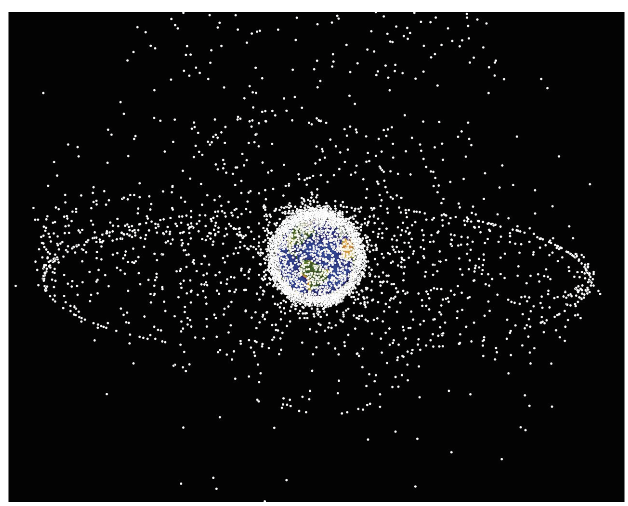

Figure 10. Computer-generated image of objects in Earth orbit. (Image from NASA Orbital Debris Program Office photo gallery.)

10-day triple-site collection campaign (October 20–29, 2009) using three geographically separated Raven-class optical sensors (see, e.g., Ref. 35) at Albuquerque, New Mexico; Maui, Hawaii; and Kwajalein Atoll, Marshall Islands. The campaign was the first historical instance of long-baseline simultaneous collections at three sites by Raven-class sensors on common objects. The project developed an optical image screener that detects RSO with weak signatures, extracts target features, and generates precise measurement error estimates. The screener operates with low detection thresholds to find candidates with weak signatures in the upstream data. The project also adapted the MHDA component used for ground surveillance applications to the space domain by incorporating the appropriate sensor measurement models corresponding to the Raven-class optical sensors and target dynamics models for propagating satellite motion. These UDF prototype components processed raw imagery, detected a variety of geosynchronous satellites, and fused the extracted measurement data to track the satellites with accuracy that was an order of magnitude better than that available from the Space Catalog. They also discriminated closely spaced geostationary satellites that are commonly cross-tagged by current operations (see Refs. 36 and 37).

During 2010–2011, APL applied its UDF prototype as a major component of a project sponsored by AFRL. The aim was to demonstrate proof of concept for a semi-automated dynamic sensor tasking capability with the goal of supporting Joint Space Operations Center's rapid decision making in scenarios where the current deliberative, manually intensive process for tasking the U.S. Space Surveillance Network (SSN) is insufficiently responsive. The project used UDF software components operating in a closed feedback loop with sensor resource management software components to continually retask space-observing sensors to respond quickly to urgent information needs while maintaining minimal aggregate uncertainty (one realization of the CLCISR concept described in the *Asset Employment Efficiency* section). These components were exercised within a dynamic SSN sensor tasking simulation test bed, which was used to conduct simulation-based performance assessment experiments and to quantify performance benefits in terms of track maintenance (target custody), search efficiency, and responsiveness to emergent information needs and changing priorities (see Ref. 38 for details). In 2012, APL continued this work in support of the AFRL effort on the Ibex program, which was jointly sponsored by Defense Advanced Research Projects Agency (DARPA) and Air Force Space Command.

Irregular Warfare

The U.S. DoD defines irregular warfare as "a violent struggle among state and non-state actors for legitimacy and influence over the relevant populations." It is characterized by the use of "indirect and asymmetric approaches, though it may employ the full range of military and other capabilities, in order to erode an adversary's power, influence, and will."[39] The types of operations conducted as part of irregular warfare include counterinsurgency, counterterrorism, and unconventional warfare. These operations typically require time-sensitive prosecution of hidden, fleeting, and maneuvering targets, including individual people. In response to stressing demands on persistence and timeliness of ISR, the military has fielded a large number of surveillance assets to theaters of operation, producing an overwhelming volume of imagery, FMV, and other types of sensor data. UDF technology has the potential to mitigate the enormous demand on communications and computing infrastructure and on human analysts and operators and to turn all of these data into timely, actionable intelligence.

APL recently demonstrated a prototype UDF capability at the Joint Expeditionary Force Experiment 2010 (JEFX 10).[40,41] JEFX is a Chief of Staff of the Air Force-directed series of experiments that combines live, virtual, and constructive forces to create a near-seamless warfighting environment in which to assess the ability of selected initiatives to provide needed capabilities to warfighters. JEFX 10 was focused on irregular warfare and assessed emerging technologies and capabilities against several relevant operational threads, including convoy protection, urban targeting, raid insertion, and personnel recovery. The JEFX 10 infrastructure featured an airborne network used to share information among airborne platforms and ground nodes, enabling dynamic command and control of airborne platforms and weapons. The JEFX 10 live fly experiment took place at the Nevada Test and Target Range near Las Vegas, Nevada, in April 2010.

The UDF initiative supported Navy Commander Second Fleet (COMSECFLT) Maritime Operations Center (MOC) participation in JEFX 10 with the goal of enhancing the MOC ISR tasking, processing, exploitation, and dissemination capability and overall maritime domain awareness. The UDF prototype was deployed at the COMSECFLT MOC, Naval Station Norfolk, during the period April 12–23, 2010. The specific objective was to support the COMSECFLT MOC in executing irregular warfare operational threads by providing intelligence analysts and the collection manager with a common tactical picture of concurrent operations, precise georeferenced track information from multiple sources, and the ability to exploit large volumes of data rapidly and efficiently.

The UDF prototype was specialized to consume and process FMV data collected by a team of unmanned aerial vehicles (UAVs) (illustrated in Fig. 11). The JEFX 10 operational threads used several Raven and

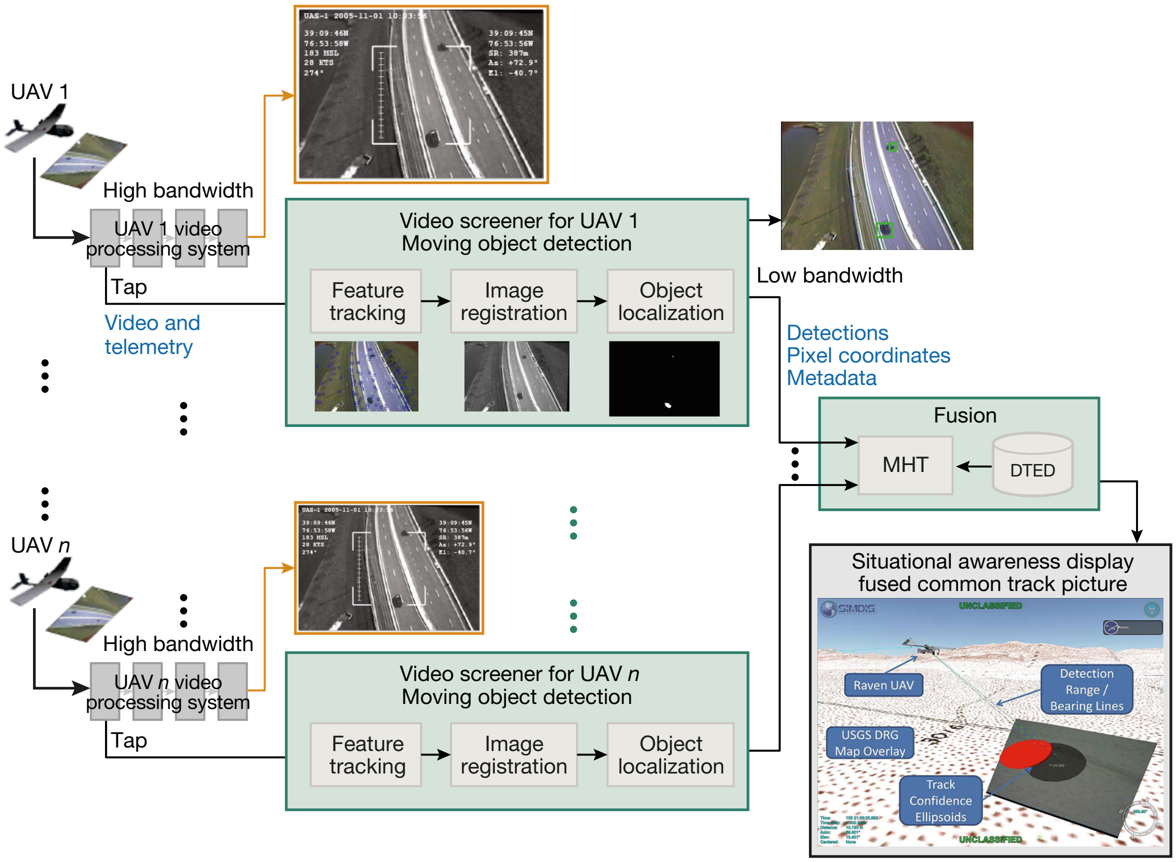

Figure 11. JEFX 10 concept for UDF of multiple UAV FMV sources. USGS, U.S. Geological Survey; DRG, digital raster graphic.

Buster UAVs as the primary FMV platforms. The prototype consisted of a video exploitation component that automatically detected moving objects (shown in Fig. 11 and by example applied to JEFX 10 in Fig. 12), a MHT that fused all of the detection data to produce a common track picture, and a display and user interface component that visualized the common track picture along with appropriate geospatial information such as maps and terrain as well as target coordinates, containment ellipses, and the source video (also shown in Fig. 11). It ran continuously and unattended during operational threads, with rare down-times requiring operator restart and very high operational availability. The system exploited a very high percentage (greater than 90%) of all CHK Raven and Buster video received at Navy COMSECFLT MOC and automatically detected moving targets, including vehicles and people. The detection and false alarm performance varied with video quality and environmental conditions. The desert background was uniform and sparse, which represented a very challenging environment for moving object detection. The UDF detection and track data were also transmitted to another prototype system developed by the Naval Postgraduate School, which used the data to recommend platform orbit adjustments to the Navy COMSECFLT MOC commanders.

More recently, in 2011, APL's Precision Engagement Business Area's Internal Research and Development (IR&D) projects extended the UDF-based FMV exploitation capability to include automated target classification using techniques developed by Murphy and colleagues and based on the highly efficient Map Seeking Circuit (MSC) approach (see the article by Murphy et al. in this issue). Automatic target classification is important

Figure 12. Automatic detection of a moving vehicle on a road in the desert during JEFX 10.

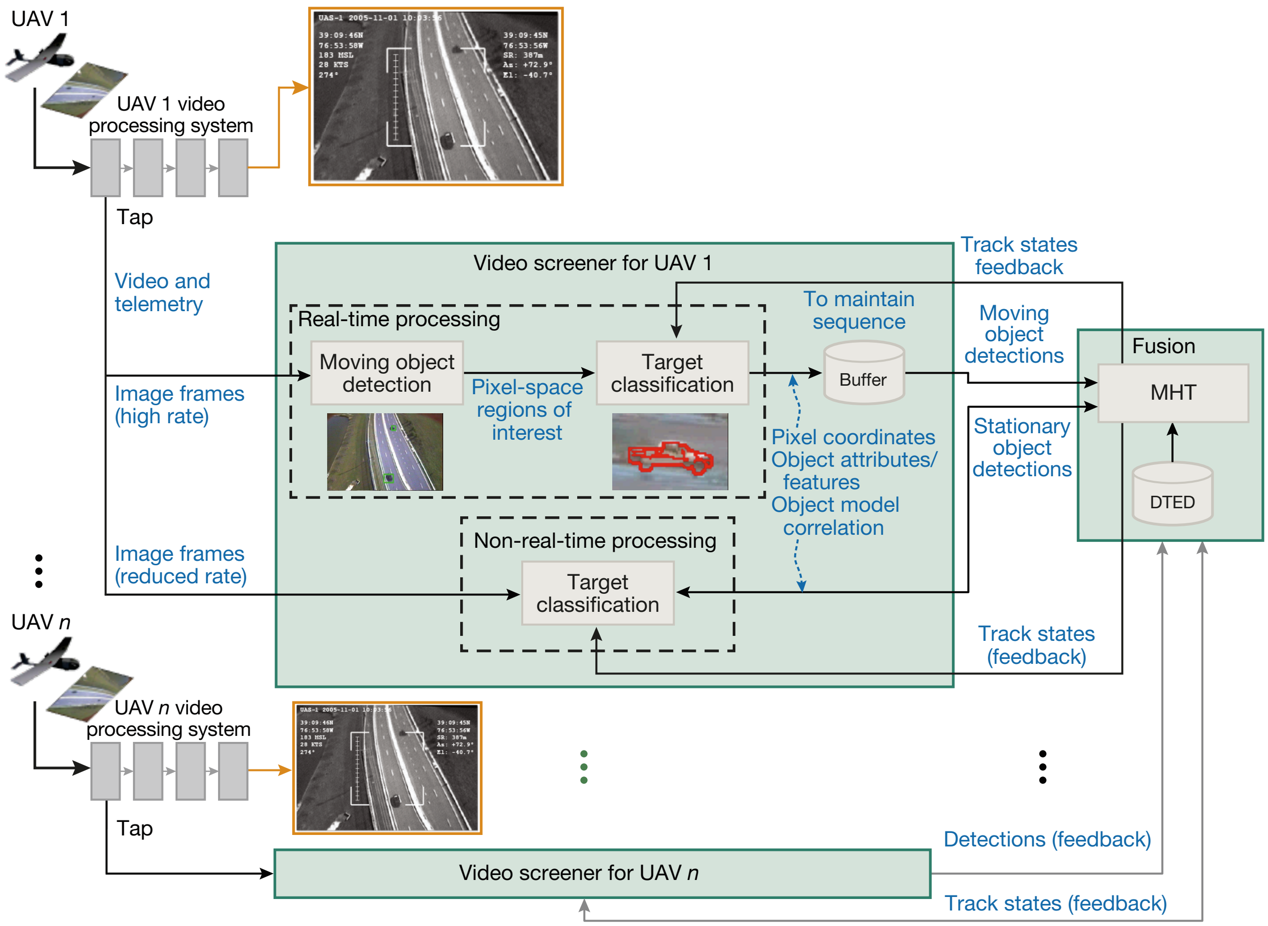

Figure 13. UDF-based FMV exploitation enhanced with automatic target classification.

for discriminating among different objects detected in the scene and correctly associating new detections with existing tracks. The idea is that the likelihood of incorrectly associating a detection with the wrong target is reduced when the classifying information can be used to strengthen the probability of candidate associations that are similar in terms of classification. The design (depicted in Fig. 13) also uses a feedback processing path that exploits the track information to enhance the target classification. Specifically, the target heading is used to reduce the search space for target orientation by the MSC algorithm.

The extended UDF prototype was tested by processing live video continuously from cameras mounted on buildings 13, 17, and 1 on APL's campus. The system successfully produced detections and tracks on moving vehicles, along with classification scores indicating the degree of confidence in classifying the object as one of the target types of interest (sedan, pickup truck, and APL shuttle bus) (shown by example in Fig. 14). The extended UDF capability supports near-real-time processing and exploitation of multiple FMV sources but is not yet mature enough to classify target types of interest with sufficient discriminating power to support fully automated target recognition. Further research and development is required. The IR&D projects also prototyped and demonstrated the feasibility of using the enhanced tracking and classification capability to transmit alerts to a mobile user equipped with an Android-based device. The user is not required to visu-

Figure 14. Example of target detections, tracks, and classification scores in processed FMV.

ally monitor the video surveillance but rather receives automated alerts on user-selected targets of interest active in a user-defined area of interest.

CONCLUSIONS

APL led a team that pioneered and matured the development of UDF techniques. UDF combines diverse sensor data tapped at points in the sensor processing chain that come well before reporting to an end user. UDF has the potential to improve the collective capability of a diverse set of sensors to detect, locate, and classify difficult targets. UDF may relax computational and communications requirements, is inherently distributed and scalable, and forms the technical foundation for a more futuristic and comprehensive concept called CLCISR. Potential benefits of UDF have been demonstrated at different levels of maturity in the laboratory and in the field for applications in time-critical targeting of relocatable ground targets, in maritime situation awareness, in space situational awareness, and in irregular warfare. Other contributions to critical ISR challenges are being investigated.

ACKNOWLEDGMENTS: The authors recognize the vast and varied contributions of the following individuals to the development of UDF applied to critical challenges in the ground, maritime, and space domains (organizations are listed in order of number of contributors; contributors within an organization are listed in alphabetical order): APL (and its consultants): Marshall H. Alworth, Randall M. Arvizu, Drew Asher, Bradley O. Bates, Chris W. Baumgart, David Bazell, David M. Berstein, Christopher M. Boswell, N. Ray Briscoe, Myron Z. Brown, Thomas H. (Buck) Buchanan, Louis P. Butler, James P. Christ, Cash J. Costello, Michael A. Cramer, Thomas B. Criss, Colleen A. Davis, Jonathan T. DeSena, Vince Doran, Thomas E. Duerr, Lien T. Duong, Daniel A. Dutrow, Lance E. Earwood, John Florence, Michael Francisco, Barry E. Fridling, Chip Frock, Lee Ann D. Fults, William J. Geckle, Jeffrey M. Gilbert, Michael T. Grabbe, Brian E. Gray, Sean D. Healy, Daniel Hein, Gina R. Hoffman-DeLeone, Tom Jerardi, Bryan A. Johnson, Jeffrey D. Jordan, Eliezer G. Kahn, Eric M. Klatt, Susan L. Kirk, Chrisanthe P. Kostalas, Larry J. Levy, Glen A. Long, Mark D. LoPresto, Kevin M. Lyons, Brooke A. Mays, Jeromy W. McDerment, Nishant L. Mehta, William A. Menner, Christopher Michaelis, Robert J. Mitchell, Patricia K. Murphy, Kathleen A. Newtson, Gregory A. Niswonger, Brian C. Page, Cameron K. Peterson, Keith S. Peyton, David W. Porter, Vignesh R. (Vinny) Ramachandran, Roy Robinette, Pedro A. Rodriguez, Donald A. Ruffle, John Samsundar, Dana M. Schlichting, David P. Silberberg, James C. Spall, Tom Spisz, Jacqueline A. Stock, Lora S. Vanegas, Meta M. Voelker, Bill Walker, William C. Walton, Adam S. Watkins, Frank B. Weiskopf, Michael J. White, Norbert White, Stephen E. Wieprecht, Michael P. Williams, Melissa K. Wong, and James C. Wright; Northrop Grumman Corporation: Tyler Elliot, Kyle Gordon, John Noone, Ras Smith, and Chris Tyler; Atlantic Aerospace Corporation: Pete Kittredge, Tamar Peli, and Mon Young; Charles S. Draper Laboratory: Frank Geisel, Tony McPherson, and Judy Miller; General Dynamics/Veridian: Mark Carlotto, Marty Liggins, and Mark Nebrich; Lockheed Martin Corporation: Chung-Fu Chang, Don Sopiak, Pat Spicer, and Joe Swartz; Orincon Corporation: James Hawker, Dennis Likes, and Peter Shea; Scitor Corporation: Ed Elguthon, Bob Halcomb, and Mike Rector; Dynetics Corporation: Jason Gallaspy and William Lummus; Georgia Tech Research Institute: David Millard and Russell Roberts; SAIC: Steve Sayre and Warren Smith; and Zel Technologies, LLC: Mike Shoop. There have been so many contributors that we sincerely hope we have not missed anyone.

REFERENCES

[1]Hall, D. L., and Llinas, J., "Multisensor Data Fusion," Chap. 1, *Handbook of Multisensor Data Fusion: Theory and Practice*, 2nd Ed., M. E. Liggins, D. L. Hall, and J. Llinas (eds.), CRC Press, Boca Raton, FL (2009).

[2]Bath, W., Boswell, C., Sommerer, S., and Wang, I.-J., "Detection Systems Information Fusion," in *Johns Hopkins APL Tech. Dig.* **26**(4), 306–314 (2005).

[3]*Intelligence, Surveillance, and Reconnaissance: Preliminary Observations on DOD's Approach to Managing Requirements for New Systems, Existing Assets, and Systems Development*, GAO Report GAO-07-596T (Apr 2007).

[4]*Unmanned Aircraft Systems: Advanced Coordination and Increased Visibility Needed to Optimize Capabilities*, GAO Report GAO-07-836 (Jul 2007).

[5]*Intelligence, Surveillance, and Reconnaissance: DOD Can Better Assess and Integrate ISR Capabilities and Oversee Development of Future ISR Requirements*, GAO Report GAO-08-374 (March 2008).

[6]*Intelligence, Surveillance, and Reconnaissance: Actions Are Needed to Increase Integration and Efficiencies of DOD's ISR Enterprise*, GAO Report GAO-11-465 (June 2011).

[7]Deptula, D., and Francisco, M., "Air Force ISR Operations: Hunting Versus Gathering," *Air Space Power J.* **XXIV**(4), 13–17 (Winter 2010).

[8]DeGroot, M. H., *Optimal Statistical Decisions*, Wiley Classics Library, Wiley-Interscience (2004).

[9]Aerospace Command, Control, Intelligence, Surveillance, and Reconnaissance Center, "Study on Optimizing the Information Process for Warfighting," in *Proc. 2000 National Correlation Working Group Workshop*, Ft. Monroe, VA (May 2000).

[10]Schwaber, K., and Beedle, M., *Agile Software Development with Scrum*, Prentice Hall (2002).

[11]Porter, D. W., "Quantitative Data Fusion: A Distributed/Parallel Approach to Surveillance, Tracking, and Navigation Using Information Filtering," in *Proc. Fifth Joint Service Data Fusion Symp.*, Laurel, MD, pp. 87–127 (Oct 1991).

[12]Brown, M. Z., *A First Look at SAR ATR for the Five-Minute War*, Technical Memorandum VNT-1-01U-003, JHU/APL, Laurel, MD (4 May 2001).

[13]Reid. D., "An Algorithm for Tracking Multiple Targets," *IEEE Trans. Automatic Control*, **24**(6), 843–854 (1979).

[14]Blackman, S., "Multiple Hypothesis Tracking for Multiple Target Tracking," *IEEE Aero. Electron. Sys.*, **19**(1) (Jan 2004).

[15]Bar Shalom, Y., Blackman, S., and Fitzgerald, R., "Dimensionless Score Function for Multiple Hypothesis Tracking," *IEEE Trans. Aerospace Electron. Sys.*, **43**(1) (Jan 2007).

[16]Kurien, T., "Issues in the Design of Practical Multitarget Tracking Algorithms," *Multitarget-Multisensor Tracking: Advanced Applications*, Y. Bar-Shalom (ed.), Artech House, Norwood, MA, pp. 43–83 (1990).

[17]Wigren, T., "Target-Type Probability Combining Algorithms for Multisensor Tracking," in *Signal Processing, Sensor Fusion, and Target Recognition X*, Proc. SPIE Vol. 4380, I. Kadar (ed.), SPIE, Bellingham, WA, pp. 46–62 (2001).

[18]Bar-Shalom, Y., Kirubarajan, T., and Gokberk, C., "Tracking with Classification-Aided Multiframe Data Association," *IEEE Trans. Aerospace Electron. Sys.*, **41**(3) (Jul 2005).

[19]Bierman, G. J., *Factorization Methods for Discrete Sequential Estimation*, Academic Press (1977).

[20]Bierman, G. J., Belzer, M. R., Vandergraft, J. S., and Porter, D. W., "Maximum Likelihood Estimation Using Square Root Information Filters," *IEEE Trans. Automatic Control*, **35**(12), 1293–1298 (1990).

[21]Boswell, C. M., *Geolocation Square Root Information Filter Formulation Using Measurements of Altitude and Measurements of Time Difference of Arrival*, Technical Memorandum PSA-05-023, JHU/APL, Laurel, MD (07 Sep 2005).

[22]DeSena, J. T., *Bias Treatment under Multiple Hypothesis Data Association using Square Root Information Filters for DTCWC*, Technical Memorandum GVE-07-0039, JHU/APL, Laurel, MD (6 Jun 2007).

[23]DeSena, J. T., *Simulation-Based Study of Processing Order Dependence for IMINT-ELINT Sensor Ensembles*, Technical Memorandum GVE-08-0086, JHU/APL, Laurel, MD (9 Oct 2008).

[24]Newman, A. J., and Kittredge, P., "The Winchester Experiment—Combined Feature-Based SAR Classification and Geo-Observable Multiple Hypothesis Tracking for Improved GNCST Data Association," in *Proc. 2005 MSS National Symp. on Sensor and Data Fusion*, Monterey, CA, paper SE02 (May 2005).

[25]Newman, A., Gilbert, J., Boswell, C., Murphy, P., Florence, J., et al., "Performance Evaluation of ELINT-IMINT Upstream Sensor Fusion for JBMC2," in *Proc. 2006 MSS National Symp. on Sensor and Data Fusion*, McLean, VA, paper SC02 (Jun 2006).

[26]Mitzel, G., Boswell, C., Murphy, P., Walton, W., Newman, A., Florence, J., Gilbert, J., and Stock, J., "Multimodal Upstream Data Fusion: Theoretical Basis, Initial Results, and Future Directions," in *Proc. 2006 Missile Defense and Intelligence Community Interchange Conf.*, Monterey, CA (Aug 2006).

[27]Newman, A. J., Spall, J. C., Mitzel, G. M., Porter, D. W., Walton, W. C., et al., "DTCWC Upstream Data Fusion: Feature-Based Fusion, Sensor Bias Compensation, and Northern Edge 2006 Field Demonstration," in *Proc. 2007 MSS National Symp. on Sensor and Data Fusion,* McLean, VA, paper SE05 (Jun 2007).

[28]Defense Science Board, *Report of the Joint Defense Science Board Intelligence Science Board Task Force on Integrating Sensor-Collected Intelligence* (Nov 2008).

[29]U.S. Maritime Security Policy Coordinating Committee, *National Plan to Achieve Maritime Domain Awareness for The National Strategy for Maritime Security* (Oct 2005).

[30]Chief of Naval Operations, *Navy Maritime Domain Awareness Concept* (May 2007).

[31]Barton, C., Hobbs, S., Kaina, J., Kessler, O., Lapin, B., Tollefson, E., White, F., *Maritime Domain Awareness (MDA): Data Fusion Assessment Final Report* (Jun 2006).

[32]Li, X. R., and Jilkov, V. P., "Survey of Maneuvering Target Tracking. Part V: Multiple-Model Methods," *IEEE Aero. Electron. Sys.*, **41**(4) (Oct 2005).

[33]Dyer, P., and McReynolds, S. R., "Extension of Square Root Filtering to Include Process Noise," *J. Optimization Theory Applications*, **3**(6), 147–158 (1969).

[34]McDerment, J. W., *Crowning Ships Tracks with Upstream Data Fusion*, APL IR&D Completed Project DTIC Report (Sep 2011).

[35]Sabol, C., Kelecy, T., and Murai, M., "Geosynchronous Orbit Determination Using the High Accuracy Network Determination System (HANDS)," in *Proc. 14th AAS/AIAA Space Flight Mechanics Meeting*, Maui, HI, paper AAS 04-216 (Feb 2004).

[36]Newman, A. J., Michaelis, C. H., Klatt, E. M., Spisz, T. S., Mehta, N. L., Challa, M. S., Bazell, D., and Casanova, F., "Upstream Data Fusion for Space Situation Awareness," in *Proc. 2009 MSS National Symp. on Sensor and Data Fusion*, Nellis AFB, NV, paper SA08 (Aug 2009).

[37]Newman, A. J., Michaelis, C. H., Klatt, E. M., Spisz, T. S., Mehta, N. L., and Kahn, E. G., *Upstream Data Fusion for Space Situation Awareness, Phase 2 Final Technical Report*, APL Technical Report (Jun 2010).

[38]Newman, A. J., Klatt, E. M., Martin, S., Clarke, J., Rodriguez, B., and Mehta, N., "Dynamic SSN Sensor Tasking for Space Situation Awareness," in *Proc. 2011 MSS National Symp. on Sensor and Data Fusion*, Washington, DC, paper NA04 (Oct 2011).

[39]U.S. Department of Defense, *Irregular Warfare (IW) Joint Operating Concept (JOC), Version 1.0* (Sep 2007).

[40]Funk, B., Newman, A., and Hiatt, D., "JEFX-10: Autonomy and Data Fusion," in *Proc. 2010 JHU/APL Small UAS Symp.*, Laurel, MD (Jun 2010).

[41]Funk. B., Castelli, J., Watkins, A., McCubbin, C., Marshall, S., et al., "JEFX 10 Demonstration of Cooperative Hunter Killer UAS and Upstream Data Fusion," *Unmanned Systems Technology XIII*, Proc. SPIE Vol. 8045, D. W. Gage, C. M. Shoemaker, R. E. Karlsen, and G. R. Gerhart (eds.), SPIE, Bellingham, WA (Apr 2011).

The Authors

Andrew J. Newman is a member of the APL Principal Professional Staff and Section Supervisor of the Data Fusion Analysis Section of the Weapon and Targeting Systems Group. His work at APL has focused on sensor and data fusion, target tracking, and dynamic platform and sensor resource management applied to ground, space, and maritime surveillance. He has led tasks on projects mentioned in this article, including DTCWC, UDF for SSA, and JEFX 10 UDF of UAV-Borne Video Sensors. **Glenn E. Mitzel** is a member of the APL Principal Professional Staff and Chief Scientist for the Precision Engagement Business Area. His interests include optimal estimation and control, remote sensing, and rapid prototyping. Since 2000 he has led large multidisciplinary teams developing, applying, and expanding the UDF concept. For further information on the work reported here, contact Andrew Newman. His e-mail address is andrew.newman@jhuapl.edu.

The *Johns Hopkins APL Technical Digest* can be accessed electronically at **www.jhuapl.edu/techdigest**.