

MULTI-LEVEL SENSOR FUSION ALGORITHM APPROACH FOR BMD INTERCEPTOR APPLICATIONS

Doug Allen, Bart Smith, Norman Morris, Charles Bjork and Kim Deal-Giblin,
Nichols Research Corporation, 4040 South Memorial Parkway, Huntsville, AL, 35815-1502

John Rushing CSSD-TC-SS, USASMDC
P.O. Box 1500, Huntsville, AL 35807-3801

ABSTRACT

The exoatmospheric intercept regime poses a difficult technical challenge for on board interceptor discrimination, especially against mature threats that stress the seeker discrimination timeline and ability to distinguish differences between RV and decoy target characteristics. The performance may be further degraded when the observed threat characteristics are different from a priori expectations. The Discriminating Interceptor Technology Program (DITP) is being pursued by BMDO to increase the discrimination performance robustness of NMD and TMD interceptors through fabrication and testing of advanced sensor hardware concepts and advanced sensor fusion algorithms. Advanced sensor concepts include onboard LADAR in conjunction with a multi color passive IR sensor. Intelligent processing which can effectively combine sensor data from disparate sensors by selecting and using only the most beneficial sensor phenomenology data is a critical element of future exoatmospheric interceptor systems. The major goal of these processing algorithms is: to make optimal use of the multi-sensor data in good a priori conditions and to also provide robust discrimination when confronted with off-nominal or non a priori conditions. This paper summarizes the intelligent processing algorithms being developed, implemented and tested to intelligently fuse data from LADAR and passive infrared sensors at both the feature and decision levels. These intelligent algorithms employ dynamic selection of feature sets and the weighting of multiple classifier decisions for performance optimization, while minimizing on board processor memory and throughput requirements. Feature sets can be dynamically selected based on an estimate of the individual feature confidence. Target designation decisions can be improved by fusing weighted individual classifier decisions, the outputs of which contain an estimate of the confidence of the data and decisions. The confidence in the data and the decisions can be used in real time to dynamically select different sensor feature data or to request additional sensor data on specific objects that have not been confidently identified as being lethal or non-lethal. The algorithms are implemented in C within a graphical user interface framework. The baseline set of fused sensor discrimination algorithms

with intelligent processing are described in this paper. Example results are shown for simulated sensor measurement data.

Keywords: multi-sensor fusion, discrimination, classifier confidence, feature confidence, infrared, laser radar, intelligent processing.

1.0 INTRODUCTION

The Discriminating Interceptor Technology Program (DITP) is being pursued by BMDO to fabricate and test advanced sensor hardware concepts and also advanced intelligent processing and sensor fusion algorithms to increase the discrimination performance robustness of NMD and TMD interceptors. The object discrimination algorithms and intelligent processing architecture developed for DITP may also be applicable to other surveillance programs such as SBIRs, GBR, and THAAD, etc. As seen in figure 1.1, the advanced sensor concepts being developed include onboard LADAR in conjunction with a dual waveband passive IR sensor.

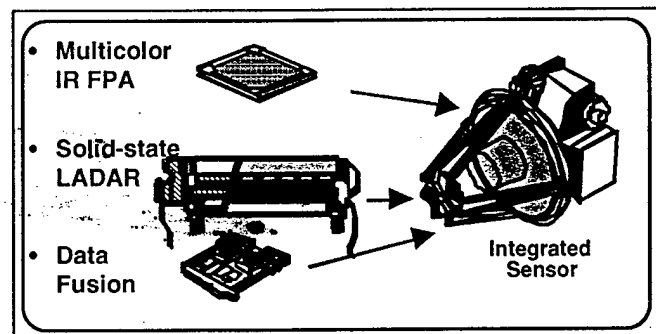


Figure 1.1 DITP Technology Demonstration

The goal of NMD and TMD missile defense systems is to protect U.S. assets from hostile threats, usually with limited interceptor inventories. Therefore, the interceptor is required to provide high quality discrimination to achieve a low RV leakage and low false alarm interceptor wastage. The exoatmospheric intercept regime poses a difficult technical challenge for discrimination against mature threats in which the offense has deliberately attempted to "match" the RV and

decoy characteristics. Far term threats are expected to be stressing to both surveillance sensors and to interceptors. DITP is being developed to provide autonomous on-board interceptor discrimination performance when provided only a cluster handover from either ground based or space based surveillance sensors.

Expected far term threats drive the interceptor IR acquisition range, resolution capability and the overall available timeline for performing discrimination. An example NMD timeline appears in Figure 1.2. Currently, NMD and TMD interceptor designs include passive only sensors. The small threat spacings of postulated threats preclude the IR sensor resolution of some objects until the engagement endgame. Even with resolution delayed, the interceptor must identify the most lethal target at any time point while maintaining the most promising targets within the sensor field of view and divert capability. Discrimination decisions are required with resolved track times as short as 10 to 20 sec, which effects performance in some cases since the sensors observe only 1/4 to 1/2 of an object's dynamic precession cycle. The addition of a LADAR sensor enhances performance in three ways: (1) provides multi-phenomenology which makes it more difficult for the offense to design penaidts that appear credible to both IR and active sensors; (2) enhances object resolution; and (3) enhances discrimination by obtaining features that extract pertinent information during short observation windows. Intelligent processing techniques, which combine sensor data from disparate sensors by selecting and using only the most beneficial sensor data, comprise the critical element of future exoatmospheric interceptor systems. Furthermore, they incorporate adaptive approaches for robust performance in off-nominal conditions using reasonable throughput and memory.

The overall intelligent processing and fusion algorithms are being developed within a simulation framework that will transition from near real-time operation on a silicon graphics workstation to a real-time implementation on an actual miniaturized interceptor processor that will be

flight demonstrated. These algorithms are first developed within the Fused Sensor Discrimination (FuSeD) simulation testbed environment, where they are supported by a graphical user interface to easily assess and evaluate algorithm and fused sensor performance. Selected algorithms will ultimately be hosted in the DITP flight fusion processor. The current baseline algorithms have been developed for implementation in a parallel processing environment to assure successful integration with the real time processor. These algorithms employ dynamic memory allocation and recursive feature extraction.

This paper emphasizes the design of the intelligent processing algorithms and focuses on the description of the multi-level fusion approach that has been implemented. Testing and evaluation of the baseline algorithms continues against both simulated and laboratory test data. Representative FuSeD evaluation results are included in this paper which highlight the utility of these fusion techniques. Through test and evaluation of the algorithms and through infusion of alternate algorithms and approaches, the baseline algorithm set will continue to be refined through the DITP flight test program.

This paper is based upon previous activities performed in support of the United States Army Space and Missile Command (USASMDC) in the development and implementation of an intelligent fusion process^[1, 2, 3, 4, 5]. The paper presents the current baseline architecture and provides example results. An overview of the intelligent processing framework and the additions made to the traditional discrimination approach are described in Section 2 of this report. In Section 3 the multi-sensor fusion approach is described, with focus on the hybrid feature set fusion approach and options for multi-classifier fusion. Section 4 presents results that show the benefit of fusion at both the feature and decision levels as the interceptor engagement unfolds.

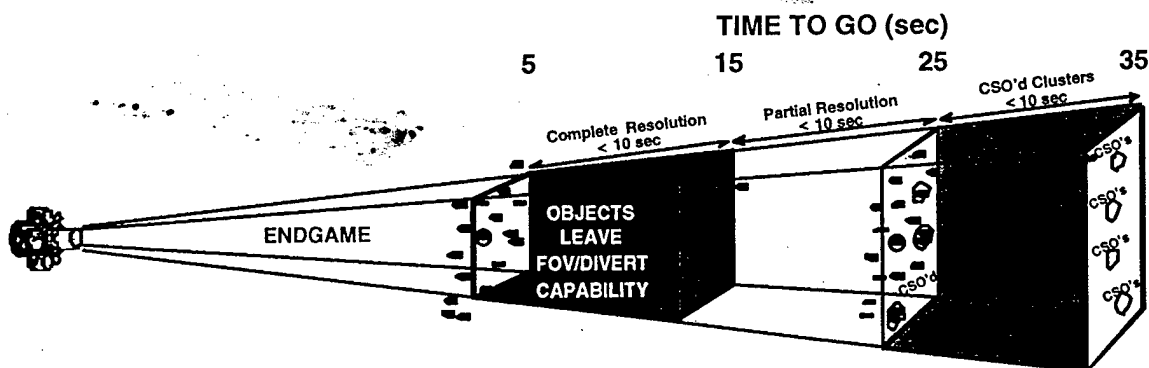


Figure 1.2 Example NMD Interceptor Timeline

2.0 INTELLIGENT PROCESSING ARCHITECTURE

An adaptive architecture framework utilizing intelligent processing (IP) techniques has been implemented to: (1) optimally use and combine multi-sensor data; (2) provide a mechanism for graceful degradation in off-nominal conditions; (3) provide feedback to the sensor resource manager via sensor prioritization requests. Further details on the IP framework exist in references [4] and [5]. The baseline fusion architecture has been implemented on the Fused Sensor Discrimination (FuSeD) Testbed and as indicated in Figure 2.1 uses the traditional discrimination process as a foundation, for extending algorithm capabilities using IP techniques.

The traditional discrimination algorithm design requires pertinent sensor information collected from a single sensor, extracted as data "features." These features are then compared to a single a priori classifier model database for target class identification (e.g., "lethal," "nonlethal" or "unknown"). Designation of the most lethal target is then made using target class probability data. Target class probabilities drive the sensor resource manager for management of sensor field of view and divert decisions. Selection of the features and the individual classifier are chosen a priori, based on analyses that define optimum performance when no "gross" mismatch between the measured data and classifier database exists.

The IP framework follows the traditional design thread and process: extraction of the pertinent sensor information through features, comparison of feature data to a priori model databases, designation of the most lethal target and output of target class probabilities to the sensor resource manager for the optimal use of sensor resources. The functions of these common modules have been extended to: (1) use inputs from multiple sensors; (2) dynamically quantify and select the optimum feature sets and weighted classifier decisions; (3) quantify the confidence in the target class probability estimate; (4) provide specific feature data requests to the multi-sensor resource manager.

Feature "sets" can be dynamically selected using estimates of feature confidence. These are determined from feature quality and weighting terms derived from the quality of sensor data and expected phenomenology. Multiple classifiers are employed which use both knowledge-based (for good a priori conditions) and adaptive clustering approaches (for off-nominal conditions) to fuse the sensor data and to provide a target lethality estimate. Target designation decisions can be made by fusing weighted individual classifier decisions whose output contains an estimate of the confidence in the data and discrimination decisions. The confidence in the data and decisions is used in real time to: dynamically select different sensor feature data; identify that off-nominal conditions exist; and to request additional sensor data on specific objects that have not been confidently identified as "lethal" or "non-lethal."

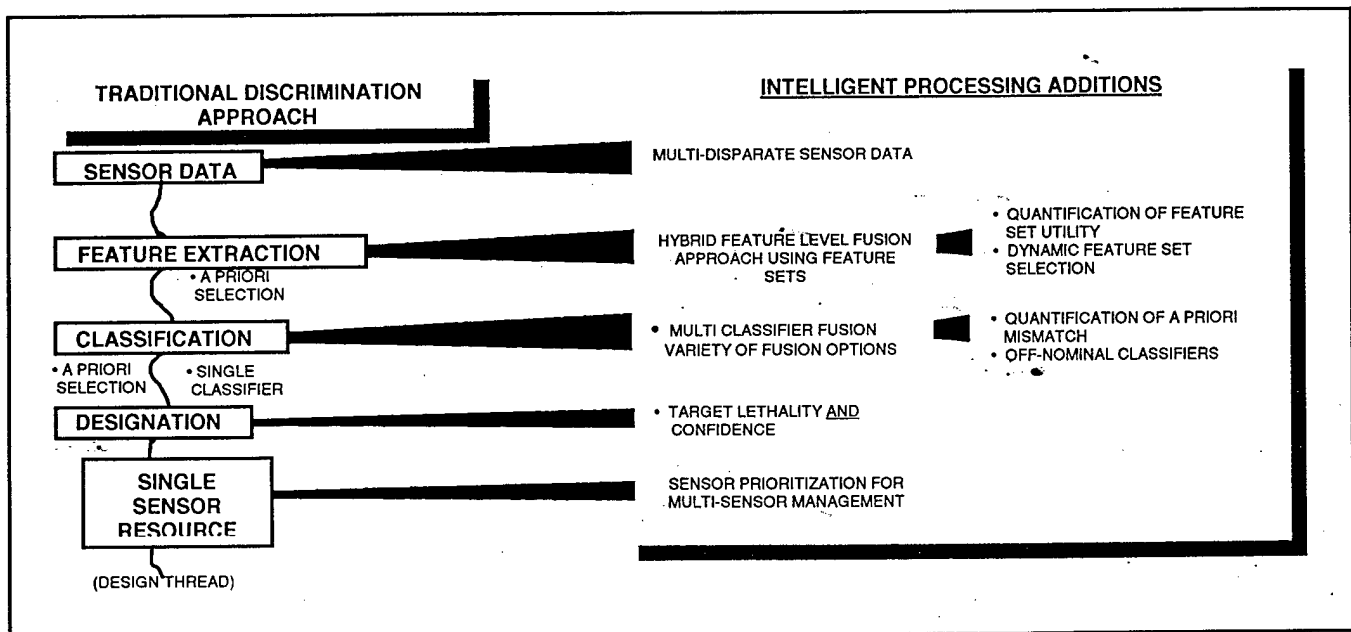


Figure 2.1 Intelligent Processing Additions to Traditional Discrimination Approach

2.1 Key Drivers on Algorithm Development

The algorithms must be implemented within the constraints of the overall fusion processor and must operate within the concepts of operation of the interceptor subsystem. The following are key drivers on algorithm requirements: (1) real time operation through the parallel processing of 4 functional processes within 400 kb of memory per process; (2) operation with large volumes of multi-sensor data (e.g., high IR updates, 3D LADAR images) that may be asynchronous and irregularly sampled; (3) provide discrimination decisions with small observation windows against advanced countermeasures and (4) provide robust performance in off-nominal conditions.

The IP architecture has been designed to meet these requirements. The intelligent algorithms have been implemented in C within a parallel messaging environment for real time operation. The algorithms are adaptive in nature and extract features utilizing a recursive approach. The algorithms employ multi-level fusion at the feature and classifier decision levels to exploit periodic and often uncorrelated spates of superior sensor feature sets or classifier performance, none of which may be reliably predicted throughout the entire engagement, against all scenarios or in all conditions.

3.0 MULTI-LEVEL FUSION APPROACH

The IP architecture provides for fusion of sensor data at the feature level (e.g., feature fusion) and at the classifier level (e.g., classifier or decision fusion). The object of performing fusion at both the feature and decision levels is to capitalize upon the following: (1) multi sensor phenomenology can benefit performance; (2) no unique set of features exist that provide optimum performance in all conditions; (3) individual sensor data quality can vary throughout an engagement and (4) the actual threat or sensor characteristics observed may not match a priori expectations. The object of fusion is to provide for optimal performance against a range of conditions.

The multi sensor fusion approach implemented in the baseline architecture is summarized in Figure 3.1. One or more feature vectors are formed through the mechanism of "feature sets" and an estimate of each feature set utility (e.g., confidence) is provided to the target designator. For each object and feature set, the respective classifiers determine target class probabilities and "confidence" factors. Combinatorial methods are

used to fuse these classifier decisions into weighted class probabilities. This provides for a smooth transition between good and poor a priori environments, the categorization of objects (lethal/nonlethal/unknown) and the output of combined decision confidence values to the DITP fusion processor. These parameters form the basis for sensor resource manager requests to obtain additional sensor data on specific objects. A feedback loop is also included in the algorithm design to provide information to the dynamic feature selection module when a poor a priori situation is detected. Feature sets, dynamic feature selection, classifier fusion and confidences will be discussed in further detail in the following sections. Results showing the benefit of fusion at the multiple levels will also be presented.

3.1 Hybrid Fusion at the Feature Level

A hybrid fusion approach is implemented at the feature level. The term hybrid is used because the approach implemented allows for optimal feature selection based on features made up of either individual or multi feature vectors. This compares to the classical approach of fusing individual feature vectors into a larger multi feature vector. However, in the classical approach the number of potential combinatorial feature vectors can become quite large. For example, 8 features from a multi color IR sensor and 8 features from an active sensor yields a possibility of 65,536 feature vectors to consider. It is acknowledged that techniques for eliminating potential combinations, e.g., "pruning techniques" could be implemented. Adoption of hybrid feature sets is intended to capture the functional performance improvement of fusing feature data while minimizing onboard memory and timing requirements.

In the discrimination process, features which are traceable to target phenomenology are extracted from the IR and LADAR sensor measurements. Differences between lethal and non-lethal class features result from differences in shape, mass distribution, surface and structural materials (both thermal and spectral), and dynamics. A priori features are characterized based on the sensor measurement capability, measured field data and projected intelligence information. However, as mentioned previously, the utility of specific features will vary; no unique set of features exist that provide optimum performance in all conditions. Individual sensor data quality can vary throughout an engagement and the actual threat or sensor characteristics observed may not match a priori expectations.

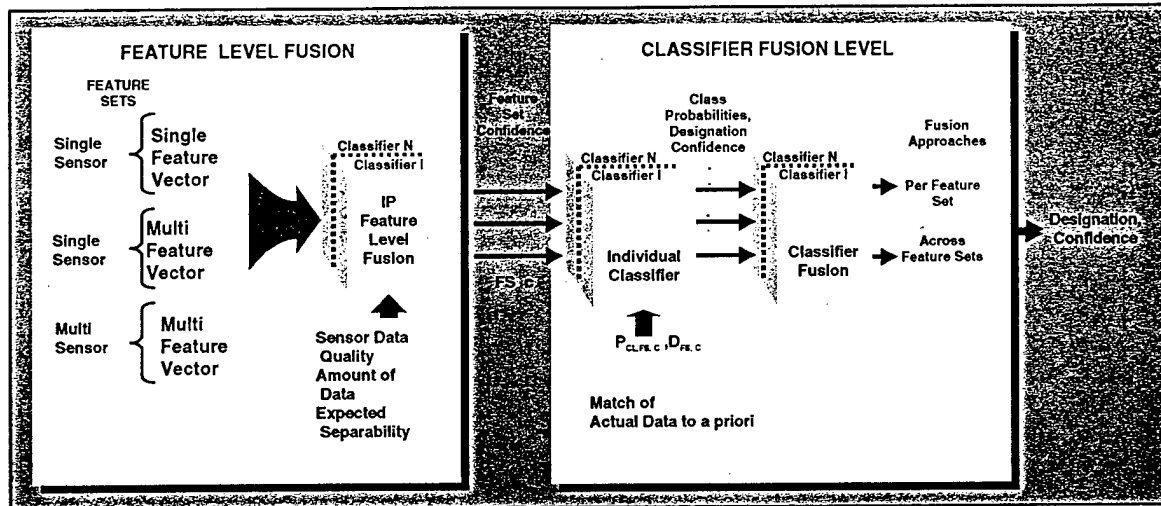


Figure 3.1 Multi-Level Fusion Architecture

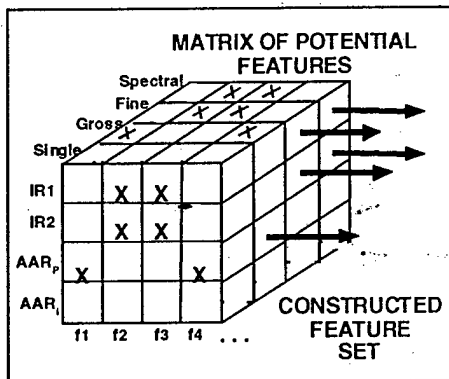


Figure 3.1.1 Feature Set Definition

Feature set composition is illustrated in Figure 3.1. A feature set can be composed of: (1) an individual sensor single feature vector (2) an individual sensor multi-feature vector and (3) a multi-sensor multi-feature vector. An ensemble of feature sets is defined prior to an engagement to exploit specific phenomenology expected based on observation time, sensor measurement characteristics and/or threat phenomenology.

The hybrid feature fusion algorithm provides weighted feature sets to the classifier algorithms throughout the engagement timeline. Feature sets can be adaptively selected to optimize discrimination potential at any given point in the engagement timeline. For each classifier, a feature confidence is determined for each object/feature set combination and is provided to the designation algorithms to serve as a figure of merit of feature set utility. Feature confidence is derived from feature quality and feature weighting terms and is defined as a measure

of the confidence in the feature utility provided for target designation. The feature quality value is determined from the quality and quantity of the sensor measurement data used in the feature calculations. The feature weighting value is determined via an "off line" evaluation of the feature's usefulness in providing discrimination, based on expected separable behavior of the phenomenology of interest. Feature confidence is determined through application of the feature weighting terms to the feature quality terms. The feature confidence equation is defined as:

$$F_{CONF} = \sum W_F \cdot F_{QUAL} \tag{1}$$

where W_F = the feature weighting, and F_{QUAL} = the feature quality estimate. The value of F_{CONF} is accumulated over a uniform sampling of the feature domain.

Features are selected dynamically throughout the interceptor engagement by selecting the feature approach (set of features) with the highest overall confidence. A threshold can be applied to downselect features in a feature set prior to providing a feature set vector to the classifier algorithms.

3.2 Fusion at the Classifier Level

The overall objective of the classifier fusion algorithm is to improve the performance and confidence in lethal or non-lethal class decision making. The baseline IP architecture has the framework that allows the flexibility of assessing several fusion options. Figure 3.2.1 highlights these options.

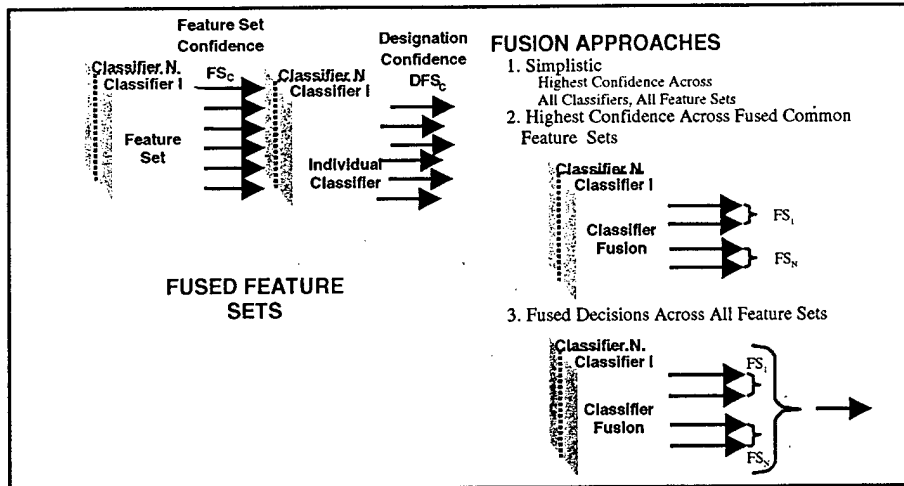


Figure 3.2.1 Classifier/Decision Fusion Options

The most simplistic option, although not recommended, is to select the most lethal object with the highest confidence across all classifiers and feature sets. This is followed with a check to see if a consensus between classifiers exists. A second option is to fuse the class probability output from all classifiers. Fusion is performed across common feature sets and is aided by use of a weighting scheme. Weighting factors are calculated according to the degree to which the observed object characteristics match the a priori expectations, and are determined for each classifier. The most lethal object with the highest confidence across all feature sets is selected. The third option combines decisions across the already fused common feature sets. These multi-classifier fusion options support optimum performance in nominal conditions. Quantification of classifier mismatch to a priori expectations allows adaptive classification to be invoked, which supports graceful degradation when confronted with off-nominal conditions.

Fusion of classifier decisions is accomplished through a weighting approach, based on the confidence in the classifier decision, and in the classifier itself. An example of the weighting term dependence on classifier design is illustrated Figure 3.2.2. The Bayesian Quadratic classifier weighting term increases with poor classifier confidence and decreases with increasing classifier confidence. In contrast the Parzen classifier weighting term increases with good classifier confidence. The strategy for this approach is to rely more heavily on the Bayesian Quadratic with poor a priori, and on the Parzen with good a priori. Note that the classifier confidence does not

affect the resulting class probabilities. The classifier confidence only impacts weighting of the decisions from each of the two classifiers.

Classifier confidence is determined for each classifier being used (Bayesian Quadratic and Parzen classifiers implemented in initial testbed release). A confidence parameter provides a measure of how well the observed phenomenology matches the expected phenomenology. Low confidence occurs when observations do not correlate well with what is expected. In this first implementation, this confidence parameter has three multiplicative constituent components: (1) Class "Similarity," (2) Class Outlier, and (3) Target Grouping, symbolized by C_{sim} , C_{out} , and C_t respectively.

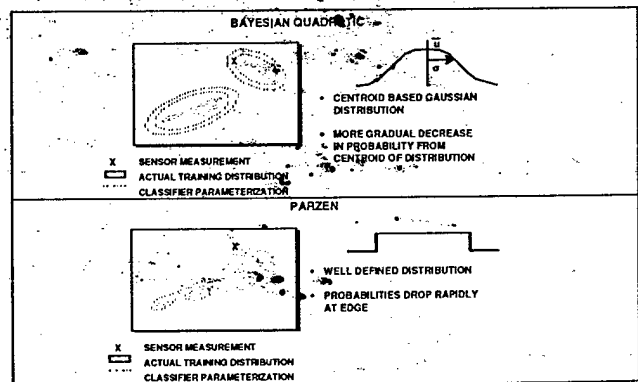


Figure 3.2.2 Classifier Characteristics

The first component of confidence, C_{sim} is illustrated in two-dimensions in Figure 3.2.3. Class membership

probabilities are used to determine the measure of confidence. The feature vectors determined for objects observed in the interceptor field of view are compared to a priori models for those objects. A weighted distance to each a priori class is computed as a first approach for confidence assessment. The next component is the outlier component, C_{out} . If the distribution of class probabilities from the sampled feature space, compared to those of a priori class probabilities, suggests that the sample is a statistical "outlier" then the confidence parameter should be appropriately reduced. The domain in which this reduction should occur is illustrated in one dimension in Figure 3.2.4. The figure shows probability densities for a single feature for two classes. The confidence reduction occurs when sample distributions occur in regions of the feature space like B, thereby being identified as "outliers." The final factor, C_c , measures the degree to which the test sample grouping of tracked objects matches the a priori grouping. The confidence is reduced if the class distribution of objects significantly differs from that of the a priori data.

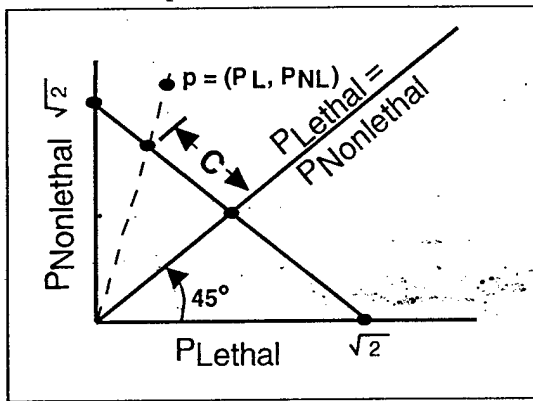


Figure 3.2.3 Similarity Component of Classifier Confidence

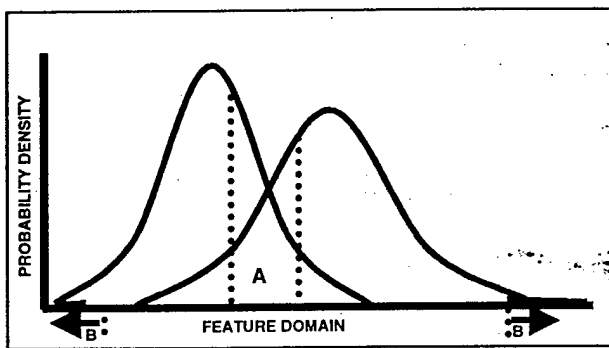


Figure 3.2.4 "Outlier" Component of Classifier Confidence

The target designation of lethal vs. non-lethal is determined through weighted probabilities and classifier confidences. Probabilities and confidences from the classifiers are fused in the following manner according to

equation (2):

$$\hat{P}_{ijk} = \frac{\sum_{M=1}^M P_{ijkm} C_{jkm}}{\sum_{M=1}^M C_{jkm}} \quad (2)$$

$$\hat{C}_{jk} = \frac{1}{M} \sum_{M=1}^M C_{jkm}$$

where k indexes objects (track files) i and j index class and feature set respectively, and m indexes the set of (m) classifiers.

The confidence-weighted class probabilities are summed over all classifiers, yielding a confidence-weighted average for each object, feature set, and class type; an average confidence is obtained by summing over all classifiers, and dividing by the number of classifiers. Thus the average confidence is available for each feature set and object, while the fused probability is available for each feature set, object, and class. Once these fused confidence-weighted class probabilities and average confidences are generated, the probabilities associated with the most lethal object are extracted, either by selecting the largest class probability from the most lethal set, or the sum of the class probabilities taken over this most lethal set. Both methods are being assessed for effectiveness, and the choice may be retained as a decision to be made in real-time as a function of conditions or engagement time. Each yields a net class probability and confidence for each feature set and object. These are provided to the target designator a final lethality determination.

The final lethality determination is made by combining the class probabilities and confidences for each object and class. The ordered pairs of class probabilities and confidences are weighted to form weighted pairs, from which a lethality measure is computed as a magnitude in the two dimensional probability-confidence domain, as indicated in Figure 3.2.5.

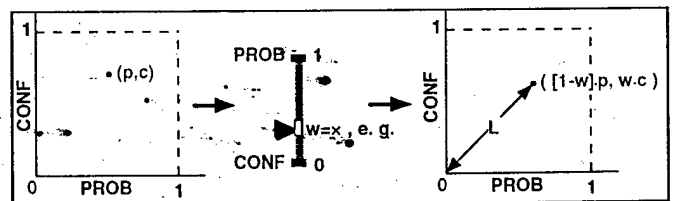


Figure 3.2.5 Designation Logic

The determination of target identification as lethal or non-lethal will be used not only for final target selection, but also in the early phases of the engagement by the field of view manager. As the engagement progresses

into midcourse and terminal phases, the lethalties of the objects will be used by the field of view manager to maintain track on the possible lethal targets for as long as possible before allowing objects to leave the field of view. This resource management is of critical importance in the successful intercept of the lethal target.

A different computed confidence value is the designation confidence. It provides a measure of the confidence in the designation decision based on both the confidence in the features and the classifier results. Computationally it is the confidence component of the chosen object with the largest lethality magnitude, averaged over all feature sets.

4.0 ANALYSIS

Results are presented that illustrate the benefit of the IP and multi fusion approach discussed in the previous sections. A stressing scenario was simulated to exercise and assess algorithm performance. The threat consisted of lethal and non lethal targets having little dynamic differences and essentially no physical or thermal mismatches. Small coning angle differences of only +/- 3 degrees existed. The close match in target characteristics make this scenario ideal to demonstrate the requirement for fused sensor discrimination due to the inability of the IR sensor alone to provide adequate performance. The engagement timeline for the scenario assumed that the IR sensor observed the threat ~ 18 seconds prior to LADAR acquisition.

4.1 Expected Benefit of Multi Sensor Feature Level Fusion

Multiple ensembles of IR and LADAR signatures are simulated to develop lethal and non lethal target distributions to be used as an a priori "learning set" for classifier association during an engagement. These distributions are used to determine the expected

Bhattacharyya separability (e.g., k-factor) and utility of various features and feature sets throughout the engagement timeline. Individual IR and LADAR features are fused together in differing combinations to create the feature sets shown in Figure 4.1.1. Feature sets FS0, FS1, and FS2 are composed of multiple LADAR features that are independent between sets. Feature sets FS3, and FS4 contain fused IR/LADAR feature vectors.

Feature Set	IR Features	LADAR Features
FS0	none	f1,f2,f3,f4
FS1	none	f5,f6
FS2	none	f7,f8
FS3	IR:f1,f2	f9,f10
FS4	IR:f1,f2	f5,f6

Figure 4.1.1 Feature Set Construction

Figure 4.1.2 illustrates the a priori feature and feature set separability across the engagement time. The results indicate in 4.1.2a that individual IR features alone provide essentially no ability to distinguish the differences between target classes. As seen in 4.1.2b, several individual LADAR features provide increased separability compared to the IR sensor.

The data supports two assertions: (1) Fused features provide increased separability and (2) the utility of features varies through the engagement timeline. The increased separability obtained by fusing features is shown in Figure 4.1.2c. For ease of comparison, Figure 4.1.3 provides the individual feature separabilities (from Figures 4.1.2a and 4.1.2b) compared to the fused feature separability at a given engagement time when IR and LADAR have observed the threat for 28 and 10 seconds, respectively.

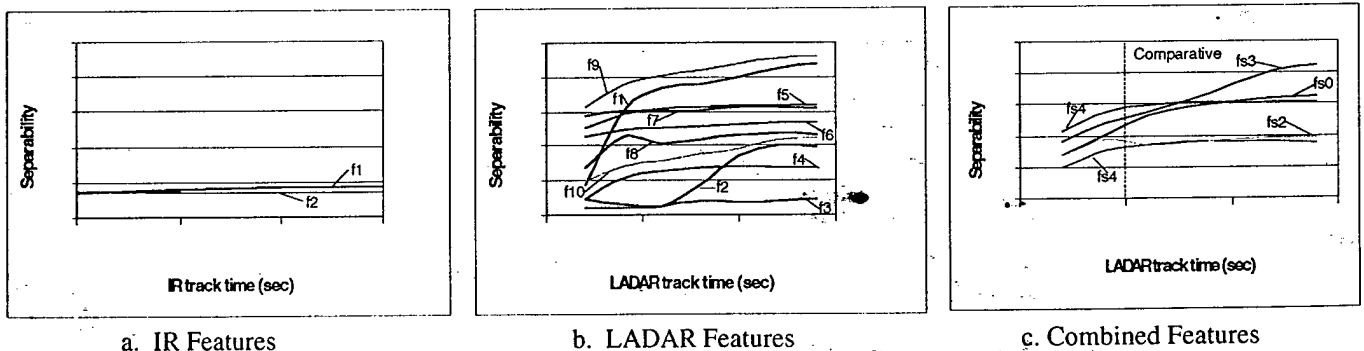


Figure 4.1.2 Individual and Fused Feature Separability

Feature Set	Indiv. Feature	Indiv. Feature Seperability	Fused Seperability
FS0	f1	1.75	2.2
	f2	0.10	
	f3	0.15	
	f4	0.6	
FS1	f5	1.5	1.77
	F6	1.25	
FS2	f7	0.60	1.8
	F8	1.2	
FS3	IR:f1	0.35	2.3
	IR:f2	0.35	
	f9	1.9	
FS4	IR:f1	0.35	2.45
	IR f2	0.35	
	f5	0.50	
	f6	1.25	

Figure 4.1.3 Comparison of Separability

Observe in Figure 4.1.2c that the utility of specific feature sets changes during the engagement. The IR/LADAR feature set FS4 provides the highest separability during the initial few seconds of LADAR track. Afterwards, the IR/LADAR feature set FS3 becomes the dominant feature set.

4.2 Classifier Level Fusion

The classifier confidence per feature set was determined for both the Bayesian and Parzen classifiers and is shown in Figure 4.2. The left column of graphs represents the Bayesian classifier results and the right column represents the Parzen classifier results. There are five feature sets, FS0 through FS4, and for each feature set there is a graph showing, as a function of time, the classifier confidences for two objects, the RV (L) and the decoy (NL). The y-axis ranges from -1.0 to +1.0, with +1.0 meaning the object appears more likely to be a lethal object and with -1.0 meaning the object appears more likely to be a non-lethal object. For these graphs, useful confidence values occur in the approximate range interval of 25 seconds. This is because LADAR data is unavailable for the first 18 seconds after initial IR data acquisition.

At each time instant, the following four cases are possible:

1. Ideal case: The lethal object classifier confidence value is close to +1.0 and non lethal object confidence value is close to -1.0.
2. Nominal case: the lethal object classifier confidence value is greater than zero and tends toward +1.0 over time while the non-lethal classifier confidence value is less than zero and tends toward -1.0.
3. Gracefully off-nominal case: the lethal object classifier confidence value is greater than zero and

4. Purely off-nominal case: Any other combination of lethal and non-lethal classifier confidence values.

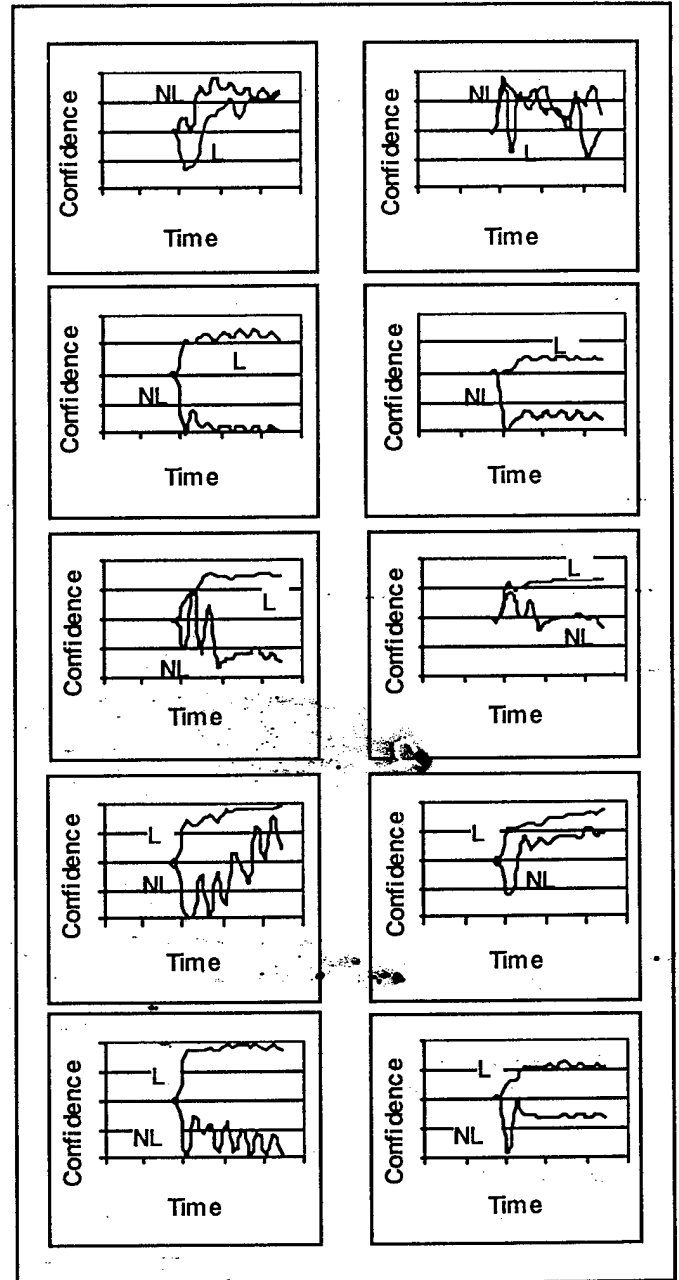


Figure 4.2 Designation Confidence

In Table 4.1 we tabulate these cases with respect to the graphs in Figure 4.2.

Table 4.1 Classifier Confidence Assessment

Feature Sets		Ideal		Nominal		Gracefully Off-Nominal		Purely Off-Nominal	
		BQC	PRZ	BQC	PRZ	BQC	PRZ	BQC	PRZ
		FS0					X		
FS1			X	X					
FS2			X	X					
FS3					X	X			
FS4			X	X					

According to this tabulation, both the Bayesian and Parzen classifiers with feature sets 1, 2, and 4 are somewhat nominal and, assuming the confidence values are truly meaningful, should predictably yield good designation results. The Bayesian classifier with feature sets 0 and 3 and the Parzen classifier with feature set 3 depict gracefully off-nominal behavior and, depending upon the behavior of the respective class probabilities, may or may not yield good designation results. The general feeling is that such cases tend toward good rather than poor designation performance. The Parzen classifier with feature set 0 depicts, for most of the engagement, purely off-nominal behavior and should predictably produce poor designation performance. Note, however, for this case the divergence that is evident toward the end of the engagement; this suggests ultimately a tendency toward nominal behavior. Some estimated feature values, particularly IR features, improve with time and this appears to be the case here. However, notice that the Bayesian classifier for the same case seems to depict the opposite behavior. This reinforces the assertion that it is beneficial to use multiple classifiers (e.g., Bayesian and Parzen), especially for stressing scenarios of this sort. Referring once again to Figure 4.2 and considering individual classifiers only (i.e., without fusion), if we make choices based purely on the classifier confidence values depicted there, we would choose the Bayesian classifier using feature set 1 over the entire engagement as perhaps the predictably best classifier overall. With classifier fusion, on the other hand, it appears that we might choose feature set 1 initially, but then perhaps switch over to feature set 4 shortly thereafter. There is also the possibility, at each time instant, of choosing different feature set and classifier combinations

as a function of the object type; however, this adds complexity to the current discussions and is a legitimate subject for follow-on efforts. In summary, we emphasize the inherent ability, and benefit, due to the flexibility in the use of combinations of multiple classifiers and multiple feature sets, to pick and choose those combinations at each time instant which predictably yield the best overall decisions.

4.3 Analysis Summary

For stressing threat scenarios the use of one sensor type alone will not achieve the effective performance possible using combined multi-sensors. Neither the LADAR nor IR alone can have the higher degree designation that the fusion of the two sensors provide. These sensors provide only part of the discriminating capabilities when operating individually. The fusion of the individual IR and LADAR features increases the separability of the targets and provides a higher degree of confidence when using either Bayesian or Parzen classifiers.

The use of both of these classifiers is an added advantage. In some cases, as shown in section 4.3, one classifier will perform better than the other for different feature sets. The utility of having two classifiers assures the best designation of lethality for that engagement or particular feature set.

5.0 SUMMARY

The fused sensor discrimination methodology, described in this paper, forms the basis for algorithms being developed for integration into the DITP program. The primary objective of these discrimination algorithms is to provide robust discrimination performance through fusion of multiple-disparate sensor data and the use of adaptive intelligent processing techniques. The current algorithms will continue to be refined to include fuzzy logic, knowledge based and adaptive approaches.

6.0 ACKNOWLEDGMENTS

This work is supported by the United States Army Space and Missile Command (USASMDC) under Contract DASG60-97-D-0002.

7.0 REFERENCES

1. C.L. McCullough, B.V. Dasarthy, and P.C. Lindberg, "Multi-level Sensor Fusion for Improved Target Discrimination", IEEE Conference on Decision and Control, pp. 3674-3675 Kobe, Japan, December 1996
2. K.A. Byrd, C.L. McCullough, C.A. Bjork, G. Grider, B.V. Dasarthy, B. Smith, N. Morris, "Intelligent Fusion Processing in BMD Applications", 1996 Proceeding of SPIE Sensor Fusion in Architecture Algorithm and Applications, Vol 3067, pp. 190-196, April 1997

3. K. Byrd, B. Smith, N. Morris, J. Rushing, "Intelligent Discrimination Fusion Algorithms For Interceptor Applications," AIAA/BMDO Technology Readiness Conference, San Diego, Ca, August 1997, Paper No. 05-01
4. K. Byrd, B. Smith, D. Allen, N. Morris, C. Bjork, K. Deal-Giblin, J. Rushing, "Intelligent Processing Techniques for Sensor Fusion," 1998 Proceeding of SPIE Aerosense: Sensor Fusion in Architecture Algorithm and Applications II, Vol. 3376, pp. 2-15, April 16-17,1998.
5. C. Bjork, N. Morris, B. Smith, D. Allen, J. Rushing, "Early Evaluation of Intelligent Interceptor Seeker," Proc. of 1998 IEEE Joint Conference On the Science and Technology of Intelligent Systems, NIST, Gaithersburg, MD, Sept 14-17, 1998

Further author information -

B.S.: Email: smithb@nichols.com; Telephone 256-885-7885; Fax: 256-922-8840
K.B.: Email: byrdk@nichols.com; Telephone 256-885-7831; Fax: 256-922-8840
D.A.: Email: allend@nichols.com; Telephone 256-885-7883; Fax: 256-922-8840
N.M.: Email: morrisn@nichols.com; Telephone 256-885-7886; Fax: 256-922-8840
C.B.: Email: bjorkc@nichols.com; Telephone 256-885-7885; Fax: 256-922-8840
K.DG.: Email: dealgibk@nichols.com; Telephone 256-885-7885; Fax: 256-922-8840
J.R.(correspondence): RushingJ@smdc.army.mil ; Telephone 256-955-3876; Fax: 256-955-4338

PLEASE CHECK THE APPROPRIATE BLOCK BELOW:

-AO # _____

_____ copies are being forwarded. Indicate whether Statement A, B, C, D, E, F, or X applies.

DISTRIBUTION STATEMENT A:
APPROVED FOR PUBLIC RELEASE: DISTRIBUTION IS UNLIMITED

DISTRIBUTION STATEMENT B:
DISTRIBUTION AUTHORIZED TO U.S. GOVERNMENT AGENCIES ONLY; (Indicate Reason and Date). OTHER REQUESTS FOR THIS DOCUMENT SHALL BE REFERRED TO (Indicate Controlling DoD Office).

DISTRIBUTION STATEMENT C:
DISTRIBUTION AUTHORIZED TO U.S. GOVERNMENT AGENCIES AND THEIR CONTRACTORS; (Indicate Reason and Date). OTHER REQUESTS FOR THIS DOCUMENT SHALL BE REFERRED TO (Indicate Controlling DoD Office).

DISTRIBUTION STATEMENT D:
DISTRIBUTION AUTHORIZED TO DoD AND U.S. DoD CONTRACTORS ONLY; (Indicate Reason and Date). OTHER REQUESTS SHALL BE REFERRED TO (Indicate Controlling DoD Office).

DISTRIBUTION STATEMENT E:
DISTRIBUTION AUTHORIZED TO DoD COMPONENTS ONLY; (Indicate Reason and Date). OTHER REQUESTS SHALL BE REFERRED TO (Indicate Controlling DoD Office).

DISTRIBUTION STATEMENT F:
FURTHER DISSEMINATION ONLY AS DIRECTED BY (Indicate Controlling DoD Office and Date) or HIGHER DoD AUTHORITY

DISTRIBUTION STATEMENT X:
DISTRIBUTION AUTHORIZED TO U.S. GOVERNMENT AGENCIES AND PRIVATE INDIVIDUALS OR ENTERPRISES ELIGIBLE TO OBTAIN EXPORT-CONTROLLED TECHNICAL DATA IN ACCORDANCE WITH DoD DIRECTIVE 5230.25, WITHHOLDING OF UNCLASSIFIED TECHNICAL DATA FROM PUBLIC DISCLOSURE, 6 Nov 1984 (Indicate date of determination). CONTROLLING DoD OFFICE IS (Indicate Controlling DoD Office)

This document was previously forwarded to DTIC on _____ (date) and the AD number is _____.

In accordance with provisions of DoD instructions, the document requested is not supplied because:

It will be published at a later date. (Enter approximate date, if known).

Other. (Give Reason)

DoD Directive 5230.24, "Distribution Statements on Technical Documents," 18 Mar 87, contains seven distribution statements, as described briefly above. Technical Documents must be assigned distribution statements.

Douglas K. [Signature]
Authorized Signature/Date

Doug Allen
Print or Type Name