# Parametric Multispectral Infrared Bulk Filtering for Theater Ballistic Missile Defense

*Kenneth V. Kitzman*

A fundamental challenge of Theater Ballistic Missile Defense (TBMD) design is to provide the weapon system with the capability of identifying the lethal object, usually a reentry vehicle (RV), from within the complex of objects associated with a ballistic missile flight. The threat complex can contain several individual large objects in addition to the RV. These might include a spent booster, an attitude control device, or both. If the threat has a solid-fuel rocket motor, numerous pieces of hot fuel debris may also be present. The potential number of such fragments is high enough to warrant concern about both the processing resources and the time required to sort through the debris. Thus, use of an efficient means to cull fuel debris measurements is desirable. An algorithm performing this function would reside in the signal processing chain prior to the target selection logic. Its goal would be to pare down the number of candidate lethal objects, thereby reducing the target selection task. Irradiance observations provided by an onboard infrared seeker operating simultaneously in two distinct wavelength regimes can be utilized to efficiently perform this function.

## INTRODUCTION

A critical function of a Theater Ballistic Missile Defense (TBMD) system involves identifying the lethal object, usually a reentry vehicle (RV), from within the complex of objects associated with a Theater Ballistic Missile (TBM) flight. Because of the hit-to-kill requirement of TBMD engagements, the RV must be discriminated from other classes of objects with a high degree of reliability. Since radar cannot see every object in the threat complex, an infrared (IR) sensor onboard the defending missile must perform this job. Owing to the nature of TBMD kinematics, the allowable window of time for performing this task is several seconds at most.

The threat complex can contain individual principal objects, such as a spent booster tank or an attitude control device. In addition, if a threat uses a solid-fuel rocket motor, hot fuel debris may be present. The potential number of such debris pieces is high enough to warrant concern that they could overload discrimination/classification capabilities.

This concern motivates the development of an efficient means of identifying observed hot fuel debris.

Such an algorithm would reside in the signal processing chain prior to the final target selection logic; its goal would be to identify fuel debris fragments so that they could be eliminated from further consideration as targets. Effectively prioritizing threat complex objects will mitigate the ultimate target selection task, leaving fewer credible candidate lethal objects for final consideration. While the primary goal of the algorithm is to recognize debris fragments, from the point of view of the weapon system it is even more important to retain any observation associated with the RV. A crucial aspect of algorithm design is to minimize the possibility of misidentifying the lethal object as debris, because such a situation would result in a failed engagement.

Existing seekers designed for TBMD applications operate in a single IR waveband. Such devices record the IR intensities and angular locations of the objects within the seeker field of view. To achieve quick results with minimal resource expenditure, it is advantageous to perform debris identification in "bulk" mode; that is, to utilize directly observed IR measurements rather than features that must be derived by processing accumulated data. Derived features consist of quantities such as mean intensity, standard deviation, and frequency, all of which must be computed from a time series accumulation of data. Directly observed features, in contrast, are restricted to instantaneous observations of angle location and IR intensity.

In terms of thermal properties—temperature and the product of emissivity and area (called *emissive area*)—an RV is relatively large and cool, whereas a fuel debris fragment tends to be small and hot. In a single IR waveband, their intensity measurements may not be distinct, since both size and temperature contribute to the IR signal. This overlap between small/hot and large/cool observations does not occur simultaneously in multiple IR wavebands, however. For this reason, a so-called "two-color" IR device can distinguish between observations that would overlap when viewed in a single waveband. This capability makes two-color IR a likely candidate for debris culling.

A threat complex prioritization scheme can use a Bayesian hypothesis test to distinguish between hot fuel debris on the one hand and larger, cooler principal objects on the other. Such an algorithm relies on statistical characterizations of observations likely to be recorded in the presence of various objects of a threat complex. Two-color IR observations depend on a number of different parameter values, including sensor waveband and noise characteristics, engagement geometry and range, and object thermal property values. Statistical models of these quantities can be folded into a computation of measured irradiance values likely to be observed in the presence of a particular object. These statistical irradiance descriptions, called *class-dependent feature distributions* (CDFDs), can form the basis for identifying objects that belong to the hot fuel debris object class. CDFDs can be broadly categorized as either parametric or nonparametric.

A parametric CDFD is constructed by matching a distribution of known form to an ensemble of measurements associated with a given object class. A common form of parametric CDFD is derived by computing a Gaussian distribution based on the sample statistics of a given ensemble. Nonparametric CDFDs are typically constructed numerically, either as histograms or through some process of numeric integration. Nonparametric CDFDs can provide more fidelity and precision in defining decision boundaries, but often at the expense of algorithm robustness.

A prioritization algorithm based on the use of nonparametric CDFDs was discussed in Ref. 1. The current article describes a parametric implementation of a similar algorithm. The following sections detail the computation of Gaussian CDFDs and describe their use as likelihood functions in the context of a binary hypothesis test designed to identify fuel debris objects. Operation of the filter is then illustrated with example calculations, the results of which are presented and discussed.

## CLASS-DEPENDENT FEATURE DISTRIBUTIONS

We denote by $\vec{\theta}$ the vector of signature-based features that will form the basis for making decisions about the identity of objects under observation. For an IR system, the features decided upon might be estimates of temperature $T$ and emissive area $\epsilon A$, in which case

$$\vec{\theta} = \begin{bmatrix} T \\ \epsilon A \end{bmatrix}. \tag{1}$$

Alternatively, a radar system might be designed to measure mean radar cross-section $\chi$, detect a peak cross-section $\pi$ over some interval of time, and compute a scintillation statistic $\zeta$. In this case, the feature vector would be

$$\vec{\theta} = \begin{bmatrix} \chi \\ \pi \\ \zeta \end{bmatrix}. \tag{2}$$

In a discussion of CDFDs, it is useful to speak of the various components of a threat complex as belonging to *object classes*. Statistical descriptions of relevant property values associated with the various object classes can then be viewed as reflecting in-class (as well as temporal) variations. Individual pieces of solid-fuel debris, for example, can be expected to exhibit some range of temperatures. This variability will appear as uncertainty

in a statistical model of temperature for the fuel debris object class. Let the notation $\Omega$ represent the set of $N$ possible object class distinctions. Objects that make up a TBM threat complex then fall into a small number of categories, including

- The reentry vehicle (RV) object class
- The fuel tank (FT) object class
- The attitude control module (ACM) object class
- The hot fuel debris (HF) object class

As understanding of the threat increases and as the threat itself evolves to exhibit additional capability, more classes could be added (for example, a countermeasures class). Next, we denote individual object classes within this set by $\omega_j \in \Omega$ for $j = 1, 2, \ldots, N$.

With these definitions established, the ability of an attending sensor to extract the feature vector $\vec{\theta}$ on an object that falls into class $\omega_j$ is given by the conditional probability density function

$$p(\vec{\theta}|\omega_j). \tag{3}$$

This CDFD takes into account the statistics of the object features, the sensor noise and sensitivity model, and the collection geometry. Once constructed, CDFDs corresponding to various object classes can be used to evaluate a particular feature realization in such a way as to allow inferences regarding the object class from which the realization originated. Gaussian CDFDs associated with a two-color IR sensor are discussed and mathematically derived in the following section.

## Two-Color IR CDFDs

Consider the feature vector $\vec{e}$ associated with a two-color IR sensor; it comprises $E_1$ and $E_2$, the irradiance values directly measured in the two seeker wavebands. In such a case,

$$\vec{\theta} = \vec{e} = \begin{bmatrix} E_1 \\ E_2 \end{bmatrix}. \tag{4}$$

Reference 2 derived expressions for nonparametric CDFDs on $\vec{e}$ calculated via numerical integration of an expression involving several conditionally independent statistical distributions. Later in this section, the construction of Gaussian CDFDs is outlined. Although a less numerically intensive task, this nonetheless still requires the specification of several parameters related to IR feature observation.

Measured irradiance is a function of viewing geometry, object thermal properties, and seeker noise and wavelength parameters. Object thermal properties—temperature and emissive area—can be modeled as dependent on the class to which the object belongs. Table 1 lists the complete set of parameters involved in this discussion.

## CDFD Parameter Values

Parameters that will be known during an engagement can be fixed at discrete values, while others must be varied to produce a representative dispersion of measurements. It is useful to model these variable parameters using probability distributions.

The boundaries of the two sensor spectral bands are fixed wavelength $\lambda$ values. The sensor modeled in this study operates simultaneously in two long-wave IR bands ($B_1$ and $B_2$ of Table 1), chosen to exhibit sensitivity to the cool temperatures likely to accompany an exoatmospheric RV.

Seeker operational range is allowed to vary uniformly over a 20-km span centered about a representative TBMD detection range. Object thermal properties are also allowed to vary: temperature variability is modeled using a Gaussian distribution, while emissive area is varied uniformly. This choice implies a relatively high degree of uncertainty about emissive area, which reflects the contributions of variations in surface emissivity, aspect, and solid-body motion. Mean and variance values of the temperature distribution are specified according to object class. Extrema of the uniform emissive area distribution are similarly class-specific. The current study encompasses the four object classes defined previously. Their temperature and emissive area parameter values are specified in Table 2. These representative values were chosen for illustration; they are not associated with a particular threat.

Random draws from the distributions given in Table 2 facilitate radiometric calculations. The spectral intensity of an object is computed first by integrating the Planck function corresponding to the object temperature over the spectral band of interest and then multiplying the resulting value by the object emissive area. This quantity

**Table 1. Two-color IR measurement parameters.**

| Parameter | Description |
|---|---|
| $B_1$ | First contiguous IR band $[\lambda_{11}, \lambda_{12}]$ |
| $B_2$ | Second contiguous IR band $[\lambda_{21}, \lambda_{22}]$ |
| $\sigma_1$ | Band 1 noise model |
| $\sigma_2$ | Band 2 noise model |
| $E_1$ | Measured irradiance in band 1 |
| $E_2$ | Measured irradiance in band 2 |
| $\varepsilon A$ | Object emissive area |
| $T$ | Object temperature |
| $R$ | Line-of-sight range |
| $\omega_j$ | $j$th object class |

**Table 2. Class-dependent distribution parameters for temperature and emissive area.**

| Object class | Temperature, $T$ (K) | | Emissive area, $\varepsilon A$ ($m^2$) | |
|---|---|---|---|---|
| | Mean, $\mu_T$ | Standard deviation, $\sigma_T$ | Minimum, $\varepsilon A_{min}$ | Maximum, $\varepsilon A_{max}$ |
| RV | 350 | 10 | 0.960 | 1.040 |
| HF | 1200 | 300 | 0.002 | 0.018 |
| ACM | 350 | 14 | 1.200 | 1.800 |
| FT | 450 | 45 | 2.500 | 7.500 |

Note: Temperature is modeled as a Gaussian random variable, $T \sim N(\mu_T, \sigma_T)$, while emissive area is modeled with a uniform distribution, $\varepsilon A \sim U[\varepsilon A_{min}, \varepsilon A_{max}]$.

is then divided by the square of the line-of-sight range to the object to compute irradiance arriving at the seeker. Application of a seeker noise model converts this presented irradiance value to the corresponding measured quantity. Following a model derived in Ref. 3, sensor noise behavior in the current investigation is characterized using additive and multiplicative components.

Additive noise, which corresponds to noise-equivalent irradiance (NEI), is the aggregate result of such phenomena as detector thermal noise, self-emission of the seeker optics, and post-detector electronics noise. A random effect that is added to the presented irradiance value, NEI tends to dominate detection performance. It is often listed in units of femtowatts (fW) per square centimeter (1 fW = $10^{-15}$ W). A multiplicative noise component is used to model the spatial nonuniformity of electro-optical gain terms on a detector array; this nonuniformity is caused by the inherent imprecision in the array manufacturing process. Because its effect is proportional to signal strength, gain nonuniformity is modeled using a multiplicative noise term. This term typically dominates the NEI noise component as the target signature grows after detection has taken place, e.g., during the discrimination phase. In the present analysis, additive and multiplicative noise terms are identical in bands 1 and 2; however, noise variability between bands is easily accommodated by the model. Both contributors are described as zero-mean Gaussian noise, with standard deviations as follows:

$$\text{Additive noise:} \quad \sigma_a = 20 \text{ fW/cm}^2 \qquad (5)$$
$$\text{Multiplicative noise:} \quad \sigma_m = 5\%.$$

Total noise figures in each band are then given by

$$\sigma_1^2 = \sigma_a^2 + \mu_1^2 \cdot \sigma_m^2$$
$$\sigma_2^2 = \sigma_a^2 + \mu_2^2 \cdot \sigma_m^2 . \qquad (6)$$

## Gaussian CDFD Construction

The statistical models and parameter values described in the preceding section can now be used to generate sample data from which model CDFDs can be derived. Figure 1 displays 1000 observations from each of the four object classes whose thermal properties are given in Table 2. The observations are computed according to the models of range, object thermal properties, and detector noise just described.

In Fig. 1, observations of FT class objects span a significantly larger portion of the measurement space than do observations from the remaining object classes. The thermal property parameters associated with the FT class cause a relatively wide dispersal of individual feature realizations. To show feature-space relationships among the other three object classes, Fig. 2 depicts the same data in a magnified form.

Model feature distributions can now be constructed using class-dependent observations. Sample statistics (mean vector $\vec{\mu}$ and covariance matrix $\Sigma$) associated with each class become the parameters that define the corresponding distribution according to the bivariate Gaussian expression

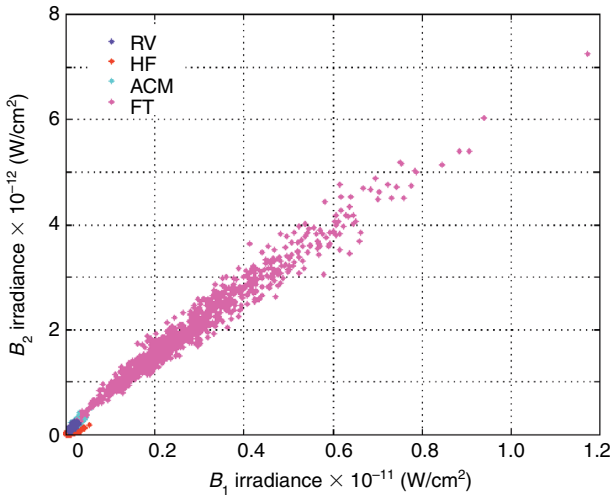


**Figure 1.** Sample observations from the four object classes. Most of the depicted measurement space is occupied by FT observations.

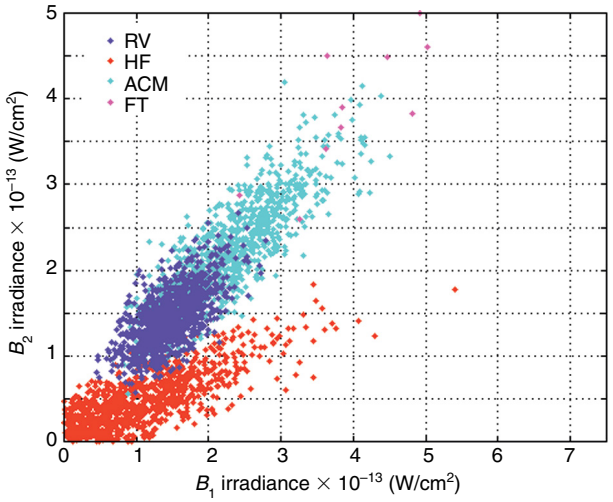


**Figure 2.** Magnified view of the sample observations in Fig. 1 with the measurement space restricted so that most observations are from the RV, HF, and ACM object classes.

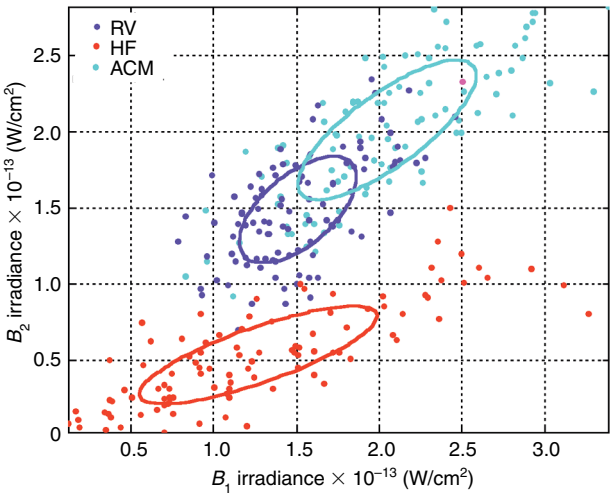


**Figure 3.** 1-σ ellipses computed from observations associated with three object classes.

$$p(\vec{e}) = \frac{1}{2\pi|\Sigma|^{\frac{1}{2}}} \exp\left[-\frac{1}{2}(\vec{e}-\vec{\mu})^T \Sigma^{-1} (\vec{e}-\vec{\mu})\right]. \quad (7)$$

Recall that $\vec{e}$ is a vector comprising irradiance values in the two sensor spectral bands,

$$\vec{e} = \begin{bmatrix} E_1 \\ E_2 \end{bmatrix}. \quad (8)$$

The variable argument of the exponential in Eq. 7,

$$(\vec{e}-\vec{\mu})^T \Sigma^{-1} (\vec{e}-\vec{\mu}), \quad (9)$$

defines a concave surface above the $E_1$–$E_2$ plane, the isocontours of which are ellipses. For example,

$$(\vec{e}-\vec{\mu})^T \Sigma^{-1} (\vec{e}-\vec{\mu}) = n \quad (10)$$

defines an ellipse centered at $(\mu_1, \mu_2)$ with ellipticity and orientation specified by the elements of covariance matrix $\Sigma$. The boundary of this ellipse encompasses the central $n$-σ portion of the corresponding Gaussian distribution. Figure 3 displays new 100-element sample observations, from which CDFDs will subsequently be constructed, for the three object classes contained in the restricted measurement space of Fig. 2. Also displayed are 1-σ ellipses derived according to Eq. 10 from respective sample statistics.

Comparing Figs. 2 and 3 reveals an example of the reduced fidelity associated with Gaussian modeling. Most of the fuel debris feature realizations in Fig. 2 are concentrated near the origin of the plot, with the frequency of observations trailing off as irradiance values increase. The symmetrical ellipse rendered according to the sample statistics of the fuel debris observations in Fig. 3 does not capture this skewness. Figure 4 depicts the joint probability distributions derived from the sample data in Fig. 3.

## BINARY HYPOTHESIS TESTING

To quickly identify solid-fuel debris objects, a binary hypothesis test (Ref. 4) can be used. In general terms, such a test is designed to process measured data and apply a *decision rule* to each observation to arrive at one of two mutually exclusive conclusions (hypotheses). The first, denoted $H_0$, is called the null hypothesis; it usually corresponds to the absence of whatever

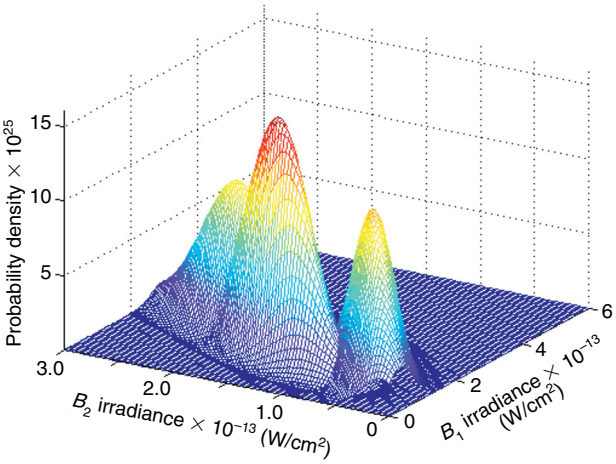


**Figure 4.** Bivariate Gaussian probability distributions based on sample data from Fig. 3.

signal, object, or condition is sought by the test. The second, $H_1$, is called the alternate hypothesis; it is associated with the presence of the particular condition under evaluation. Null and alternate hypotheses in the present context pertain to the identity of a given object under observation:

- $H_0$ indicates that the object does not belong to the HF object class.
- $H_1$ indicates that the object belongs to the HF object class.

The CDFDs described in the preceding section provide a means of processing the two-color IR measurement vector $\vec{e}$ of Eq. 8 in such a way as to select a hypothesis. For example, inserting HF object class mean and covariance values into the probability distribution given by Eq. 7 allows a computation of

$$p(\vec{e}|\mathrm{HF}),$$

or, in binary hypothesis notation,

$$p(\vec{e}|H_1). \tag{11}$$

This quantity represents the likelihood of observing a particular feature vector realization (falling anywhere on the measurement space over which the CDFD is defined), given that the object under observation belongs to the HF object class. A complementary likelihood,

$$p(\vec{e}|\sim\mathrm{HF})$$

or

$$p(\vec{e}|H_0), \tag{12}$$

conditioned this time on the null hypothesis, may be computed in similar fashion. The null hypothesis category of objects may be represented by either a single non-HF CDFD or some combination of such distributions. The point is to partition the total number of object classes into two categories, HF and non-HF (denoted ~HF), which correspond to the mutually exclusive elements of a binary hypothesis test.

The two likelihood functions defined over a measurement space region may be combined to form a *likelihood ratio* $\Lambda(\vec{e})$ given by

$$\Lambda(\vec{e}) = \frac{p(\vec{e}|H_1)}{p(\vec{e}|H_0)}. \tag{13}$$

As stated earlier, the numerator of the likelihood ratio represents the likelihood of observing a particular measurement realization $\vec{e}$, given that hypothesis $H_1$ is true, while the denominator provides the corresponding probability under hypothesis $H_0$. A likelihood ratio test is then performed by comparing the computed value of $\Lambda(\vec{e})$ to some threshold.

Associated with the two hypotheses are *a priori* probabilities $\pi_0$ and $\pi_1$, which represent an observer's information regarding which hypothesis is true before the test is conducted. Decision theory states that, on average, the probability of choosing the wrong hypothesis is minimized by employing a *Bayesian* decision criterion.[4] A Bayesian decision threshold is given by the ratio of *a priori* probabilities, weighted by the ratio of assigned costs for the two possible types of error, *false alarm* and *miss*. Error costs and *a priori* probability values must be chosen carefully to reflect the operating conditions of the algorithm. The decision threshold will vary according to changes in these values, with resulting impact on filter performance. Performance effects of various threshold values are illustrated in a discussion of receiver operating characteristics later in this article. For the present purpose of illustrating filter construction, the following results are derived based on equal error penalties and equally likely hypotheses (reflected in equal prior probabilities). In such a case, the threshold becomes unity. The test is then given by

$$\Lambda(\vec{e}) \underset{H_0}{\overset{H_1}{\gtrless}} 1. \tag{14}$$

Because the natural logarithm is a monotonic function and both sides of Eq. 14 are positive, an equivalent test may be constructed using the *log-likelihood ratio*:

$$\ln \Lambda(\vec{e}) \underset{H_0}{\overset{H_1}{\gtrless}} 0. \tag{15}$$

The threshold value of a likelihood ratio test translates into a static boundary in feature space. An observation that falls on one side of this *decision boundary* results in selection of the null hypothesis; an observation on the other side results in selection of the alternate hypothesis. The shape of this boundary depends on the probability density functions that make up the likelihood ratio, because the likelihood ratio (or log-likelihood ratio) forms a surface above the two-dimensional feature space over which the constituent distributions are defined. The selected threshold value

then corresponds to a plane above the feature space, and its intersection with the decision surface forms the decision boundary.[5]

The ~HF object likelihood function (corresponding to the null hypothesis) can consist of either a single non-HF distribution or some weighted combination of such distributions. An important aspect of processor design involves judicious representation of the ~HF object category. In the current implementation, this category is represented by the RV object class. This designation is based on the crucial importance of not misidentifying RV observations as belonging to the HF class. The CDFD in Fig. 4 corresponding to the RV object class is thus used to compute the conditional probability in the denominator of the likelihood ratio of Eq. 13.

In Fig. 2, all observations from the HF and RV classes are contained within the space bounded by $6 \times 10^{-13}$ W/cm$^2$ in the first irradiance dimension and $3 \times 10^{-13}$ W/cm$^2$ in the second. A filter designed to identify observations from the HF class thus need only be concerned with this restricted region of measurement space, and any observation that falls outside these bounds can be rejected out of hand, i.e., classified as not belonging to the HF object class. Figure 5 depicts HF and ~HF CDFDs defined over this region of measurement space.

The associated log-likelihood decision surface can now be computed according to Eq. 16:

$$\ln \Lambda(\vec{e}) = \ln p(\vec{e}\,|\,\mathrm{HF}) - \ln p(\vec{e}\,|\,\sim\mathrm{HF}) \,. \tag{16}$$

Because the two likelihood functions are given by joint Gaussian probability density functions, Eq. 15 becomes

$$(\vec{e}-\vec{\mu}_0)^T \Sigma_0^{-1} (\vec{e}-\vec{\mu}_0) - (\vec{e}-\vec{\mu}_1)^T \Sigma_1^{-1} (\vec{e}-\vec{\mu}_1) \underset{H_0}{\overset{H_1}{\gtrless}} 2 \ln \sqrt{\frac{|\Sigma_1|}{|\Sigma_0|}} \,, \tag{17}$$

where the subscripts 0 and 1 indicate RV and HF parameter values, respectively. A log-likelihood ratio decision surface is now given by the left side of Eq. 17. Note that this decision surface computes the difference between RV and HF elliptical surfaces, which are defined according to Eq. 9. The constant terms in the Gaussian likelihood functions are then combined on the right side of Eq. 17. The resulting surface appears in Fig. 6. The contour along the surface indicates intersection with the plane whose amplitude is given by the right side of Eq. 17.

## DECISION RULE IMPLEMENTATION

As filter inputs, a new set of 1000 observations associated with each object class was generated. Observations that fell within the prescribed boundaries, 0 to $6 \times 10^{-13}$ W/cm$^2$ in the first irradiance dimension and 0 to $3 \times 10^{-13}$ W/cm$^2$ in the second, appear in Fig. 7. The decision boundary associated with the specified threshold value is also displayed, clearly separating the majority of HF-class observations from those of the other classes.

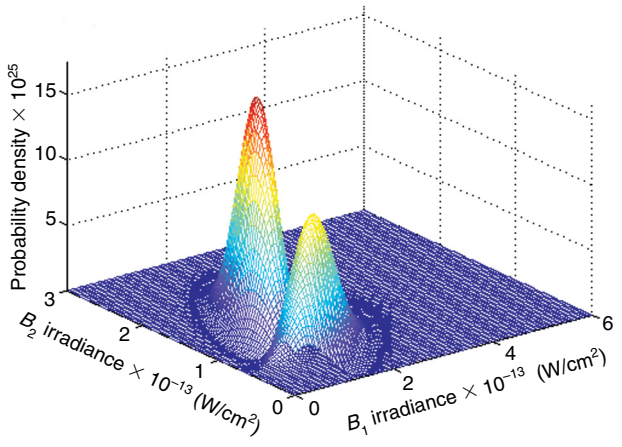


**Figure 5.** Joint Gaussian probability distributions computed for the HF and RV object classes. The RV distribution is centered roughly about $1.5 \times 10^{-13}$ W/cm$^2$ in both dimensions, while the HF distribution lies closer to the origin. The RV distribution can be used as the non-HF class (denoted ~HF) likelihood function in the likelihood ratio of Eq. 13.

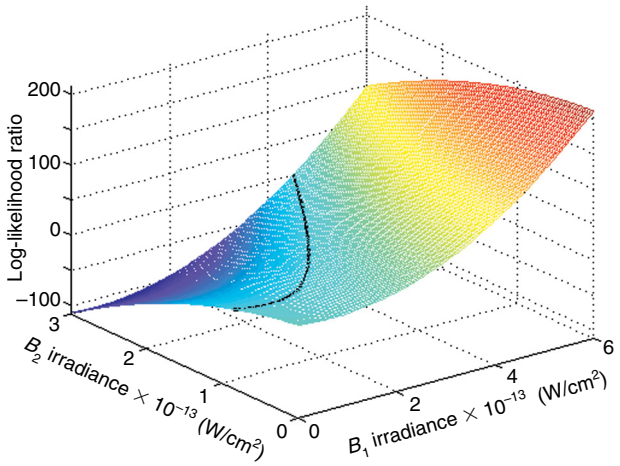


**Figure 6.** Log-likelihood ratio decision surface for binary hypothesis testing between HF and ~HF object classes. The intersection of the surface with the threshold value of $2 \ln(|\Sigma_1|/|\Sigma_0|)^{1/2}$ induces the surface contour depicted in black.

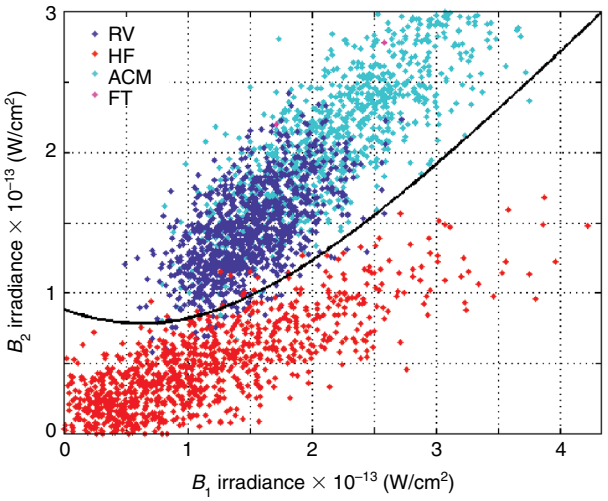


**Figure 7.** Processor inputs plotted in measurement space region restricted to contain only RV- and HF-class observations. Also plotted is the boundary resulting from a log-likelihood ratio threshold of $T = 2\ln(|\Sigma_1|/|\Sigma_0|)^{1/2}$.

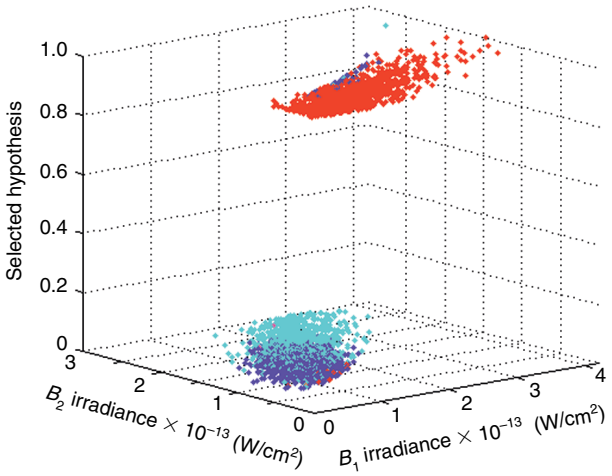


**Figure 9.** Processed data showing results of hypothesis testing. A value of 0 indicates that the null hypothesis was selected, i.e., that the observation is classified as having originated from the ~HF object classes. A value of 1 indicates the opposite, i.e., that an observation is regarded as originating from the HF object class.

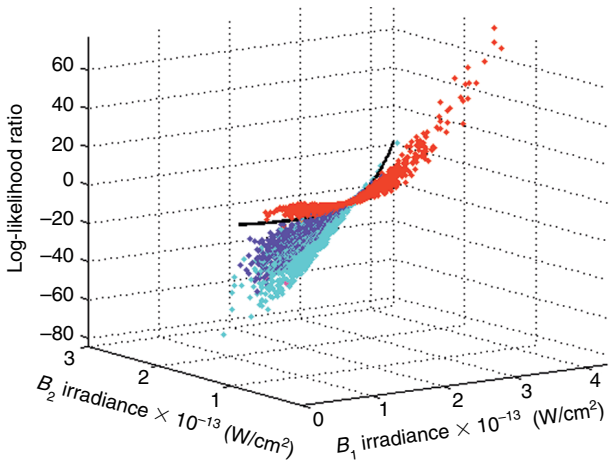


**Figure 8.** The data of Fig. 7 plotted in decision space. The decision surface value associated with each observation is plotted on the *z* axis.

Figure 8 displays the same observations in decision space. Computed for each measurement is its decision surface value, given by the left side of Eq. 17. This value appears in Fig. 8 as an amplitude in the third dimension of the plot. The decision boundary is also plotted in Fig. 8. This boundary was derived in Fig. 6 as the isocontour of the decision surface at an amplitude given by the right side of Eq. 17. Each observation whose decision surface amplitude lies below the threshold value causes the hypothesis test to conclude $H_0$, that the associated object belongs to the ~HF object category. Conversely, an observation that lies above this threshold on the decision surface results in an $H_1$ decision, and is regarded as originating from an HF-class object.

Figures 7 and 8 display observations used as inputs to the object prioritization processor. Figure 9 graphically depicts filter results. The *z* dimension amplitude of each plotted point indicates the hypothesis selected for that observation. The *x* and *y* coordinates indicate in-band irradiance values. Results are given in Table 3 as classification percentages.

## RECEIVER OPERATING CHARACTERISTICS AND EQUAL ERROR RATES

Table 3 indicates the performance levels of a particular implementation of a likelihood ratio hypothesis test in the presence of each of four different classes of threat complex objects. Such a processor is characterized by two components. The first is the log-likelihood ratio decision surface, depicted in Fig. 6. The second is the threshold value. The decision surface contains all of the available information regarding the measurement-space variability of class-dependent data. Threshold computation is considerably more arbitrary, as *a priori* probability

**Table 3. Likelihood ratio hypothesis test results.**

| Object class | Percent correct |
|---|---|
| RV | 96.0 |
| HF | 98.5 |
| ACM | 99.6 |
| FT | 100 |

Note: Percentages based on sample sizes of 1000. The notion of "correctness" corresponds to whether an observation was properly identified as belonging to the HF object class or as not belonging to this class (~HF).

and cost values are often merely educated guesses. The log-likelihood ratio function and resulting decision surface do not depend on these values, however. This invariance permits the design of a processor with a floating threshold value that can accommodate subsequent revisions.

Filter performance is often described using two parameters: *probability of detection* and *probability of false alarm*. In the present context of a processor designed to identify solid-fuel debris objects, a false alarm occurs when a non-HF-class observation results in a selection of the alternate hypothesis. In decision space, this condition translates to a likelihood ratio computed for a non-HF observation that lies above the decision threshold. Conversely, filter probability of detection corresponds to the rate at which observations originating from the HF object class cause the filter to choose the alternate hypothesis. This occurs when the likelihood ratio associated with an HF-class observation lies above the decision threshold. Probability of detection corresponds to the numerical complement of the *probability of miss*.

Recall that the filter just described was implemented utilizing a Bayesian threshold, i.e., a threshold value systematically computed based on error costs and *a priori* probabilities of each hypothesis. The probabilities of false alarm and detection associated with this threshold can be approximated by averaging the number of times the respective conditions occur. Such a process is tantamount to Monte Carlo integration. For instance, out of 1000 HF-class observations in the test sample, 985 resulted in log-likelihood ratios whose values were greater than zero. Table 3 thus reports the rate of correct classification for the HF object class, i.e., the probability of detection, as 98.5%. The corresponding false alarm rate is computed by adding the number of incorrect decisions within the three non-HF object classes and dividing by the total number of non-HF observations. From Table 3, the number of incorrect classifications for the RV class is 4% of 1000, or 40. For the ACM class, this number is 0.4% of 1000, or 4. None of the FT class observations was wrongly identified. Therefore, the probability of false alarm is 44/3000 or 1.47%.

These probabilities are associated with a single threshold value. Alternate filter implementations, employing the same decision surface but with different threshold values, will vary the level of performance. If the threshold is lowered, for instance, the probability of detection will rise, but most likely at the cost of an increased false alarm rate. Alternatively, the threshold could be raised to reduce the number of false alarms, with a corresponding decrease in detection performance. A means of summarizing the relationship between the two probabilities across the range of possible threshold values is the *receiver operating characteristic* (ROC). The ROC is a plot of filter probability of detection versus probability of false alarm for all possible threshold values.

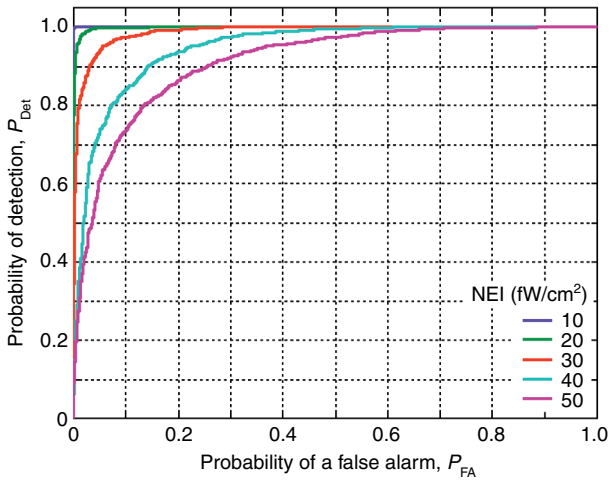


**Figure 10.** ROCs that display performance degradation with increasing NEI levels. CDFDs corresponding to various levels of additive noise were used to construct decision surfaces. One-thousand–element samples from the four object classes were then projected onto each surface and their amplitudes used to derive the curves. Each point on a given curve represents a false alarm and detection probability pair ($P_{FA}$, $P_{Det}$) associated with a particular threshold level.

Figure 10 displays ROCs associated with various levels of detector NEI. The curve labeled 20 fW/cm$^2$ is derived from the decision surface displayed in Fig. 6 and shows the relationship between probability of detection $P_{Det}$ and probability of false alarm $P_{FA}$ at all possible threshold values on that decision surface. As such, the curve passes through the ($P_{FA}$, $P_{Det}$) values (0.0147, 0.985), called the *Bayes point* of the ROC. This point consists of the false alarm and detection probability coordinates, derived in the preceding paragraphs, which correspond to the Bayes threshold. Multiple ROCs facilitate the comparison of filters derived from different decision surfaces, and thus represent a means of conducting a variety of sensitivity analyses. Additional curves in Fig. 10 display the effects on filter performance of varying the level of additive sensor noise. Better filter performance, resulting in a sharper ROC curve, is realized at lower NEI levels. Similar studies could be conducted to examine the performance impact of variations in sensor bands, multiplicative noise, range uncertainty, or any other parameter that affects likelihood ratio calculation.

The ROC is a concise graphical way to view filter performance across various threshold values. The curve can be further reduced to scalar form by employing the notion of *equal error rate*.[5] The equal error rate of a filter is found by locating the threshold value $T_{EER}$ for which the probabilities of false alarm and miss are equal. This common probability is the equal error rate. A low value of equal error rate corresponds to good processor performance; higher rates indicate decreasing functionality. The equal error rate of a likelihood ratio test corresponds to the intersection of the line $1 - P_{Det} = P_{FA}$ with the appropriate ROC. Equal error rates for the

tests that correspond to various noise levels are given in Table 4.

## CONCLUSIONS

Identification of hot fuel debris fragments can play an important role in discriminating the RV from the remaining TBM threat complex objects. Because of the short time available to perform the discrimination function, it is crucial to carry out this prioritization task quickly and efficiently. Two-color IR observations have two properties that suit them to this task: they can distinguish between small/hot objects and large/cool objects, and they are directly observable by a dual-band sensor.

Multiple sources of uncertainty exist in TBMD sensing. Such sources include measurement noise from the attendant sensor, variability in the thermal property values of the objects under observation, and uncertainty in line-of-sight ranges at which these observations will be made. The construction of Gaussian class-dependent feature distributions provides a systematic way to reflect (in measurement space) probabilistic information regarding the values of these parameters. Such feature distributions corresponding to various object classes can then serve as components in a likelihood ratio hypothesis test designed to efficiently identify hot fuel debris fragments.

**Table 4. Equal error rates for the likelihood ratio tests depicted in Fig. 10.**

| NEI ($fW/cm^2$) | Equal error rate |
|---|---|
| 10 | 0.0010 |
| 20 | 0.0220 |
| 30 | 0.0618 |
| 40 | 0.1103 |
| 50 | 0.1595 |

It should be reemphasized that the current algorithm is intended to function as a single component in the overall process of discrimination. By efficiently eliminating hot fuel debris observations from further consideration, an effective prioritization scheme can reduce the number of credible candidate threat objects from which the RV must ultimately be discriminated. Such a process must be carried out with two objectives in view: (1) quick and efficient identification of hot fuel objects and (2) retention of RV observations. Probability of error analysis facilitates evaluation of filter performance with respect to these two goals. The algorithm employed in this study achieves both objectives, given adequate signal-to-noise ratios.

## REFERENCES

[1]Kitzman, K. V., and Silberman, G. L., "Multispectral IR Bulk Filtering via Likelihood Ratio Hypothesis Testing," in *1999 National Fire Control Symp. Proc.*, Air Force Research Laboratory, Wright-Patterson Air Force Base, Dayton, OH (Aug 1999).

[2]Kitzman, K. V., and Fry, R. L., "Multispectral IR Band Selection Methodology for Target Discrimination," in *1998 National Fire Control Symp. Proc.*, Air Force Research Laboratory, Wright-Patterson Air Force Base, Dayton, OH (Aug 1998).

[3]Kitzman, K. V., *Two-Component Multiplicative Noise Sensor Model*, A1F(1)97- U-122, JHU/APL, Laurel, MD (Jul 1998).

[4]Van Trees, H. L., *Detection, Estimation, and Modulation Theory, Part I*, John Wiley & Sons, New York, pp. 19–33 (1968).

[5]Silberman, G. L., "Parametric Classification Techniques for Theater Ballistic Missile Defense," *Johns Hopkins APL Tech. Dig.* **19**(3), 322–339 (1998).

ACKNOWLEDGMENTS: R. L. Fry initially provided a theoretical Bayesian formulation that facilitates the computation of nonparametric two-color IR class-dependent feature distribution and suggested its utility in classification and filtering. A. K. Chrysostomou subsequently oversaw preliminary forays, both nonparametric and parametric, into the arena of fuel debris identification. G. L. Silberman's elucidation of decision surfaces and measurement-space boundaries added a graphical dimension to the analysis that was positively revolutionary. The author is also grateful to A. N. Courtemanche and R. A. Steinberg for their insightful review of this material.

## THE AUTHOR

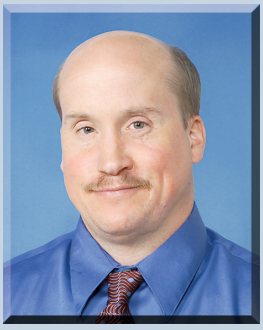

KENNETH V. KITZMAN received a B.S.E.E. in 1988 and an M.S.E.E. in 1990, both from Texas A&M University. His graduate studies focused on statistical signal processing. He joined APL in 1990 and is currently in the Electro-Optical Systems Group of the Air Defense Systems Department. He has contributed to several Navy missile programs and Theater Ballistic Missile Defense efforts in the areas of sensor modeling, feature extraction, and decision theory. His e-mail address is kenneth.kitzman@jhuapl.edu.