
Classification and Tracking of Moving Ground Vehicles

Duy H. Nguyen, John H. Kay, Bradley J. Orchard, and Robert H. Whiting

■ Airborne radars that track moving targets face the challenge of surveillance over large geographic areas where military vehicles are interspersed with civilian traffic. There is a major need to develop robust, efficient, and reliable identification and tracking techniques to identify selected targets, and to maintain tracks for selected critical targets, in dense target environments. Traditionally, tracking and identification have been considered separately; we identify a target first, and then we track it kinematically to sustain the identification. The difficulty with separate identification and tracking is that neither task is sufficient by itself to satisfy the demands of the other. We need to incorporate the distinctive target signature information into the tracker, so that identification and tracking, or signature comparison and tracking, perform together as a unit. The moving-target identification and feature-aided tracking approach described in this article combines kinematic association hypotheses with accumulated target classification information obtained from high range resolution (HRR), inverse synthetic aperture radar (ISAR), and synthetic aperture radar (SAR) signatures, to obtain improved classification and association. We describe the basics of moving-target classification and signature comparison using HRR profiles. In addition, we show improvements in HRR-based target identification using superresolution and Bayesian classification techniques. We describe how this feature information may be incorporated into kinematic tracking. Finally, we discuss methods to improve the robustness of the HRR techniques with the addition of available SAR and ISAR signatures.

CLASSIFYING AND TRACKING moving targets are difficult problems for any surveillance radar. The classification problem and the tracking problem both complement and exacerbate each other. For stationary vehicles, we can classify targets on the basis of their high-resolution synthetic aperture radar (SAR) imagery, which forms a fine-resolution two-dimensional image of the vehicle that can be compared to a previously stored template. For moving targets, however, such two-dimensional images cannot be reliably generated because of unknown vehicle motion. The primary signature that can be reliably obtained for moving targets is the high range resolution (HRR) profile. For moving targets, classification information

can be obtained by comparing an HRR profile with profiles collected from similar vehicles. However, an HRR profile contains less information about a target, compared to a two-dimensional SAR image, so reliable classification using HRR profiles relies on many looks from different aspect angles.

While the classification of moving targets relies on evidence accrued from many independent looks, such accrual can take place only if the same vehicle remains under track. Under difficult tracking situations, including low target speeds, unfavorable geometries, and dense traffic, kinematic trackers can make mistakes and associate target reports to the wrong tracks. Thus, while effective target classification could be ac-

completed with multiple-look evidence accrual, failures in track maintenance can prevent accurate target classification.

The target classification approach described in this article involves the integration of the target classification capability into a kinematic tracker, merged in such a way that the two capabilities complement and support each other. In particular, while evidence accrual for HRR-based classification relies on accurate report-to-track association, we have integrated the accumulating classification information, along with target signature information, into the association logic of the tracker. To improve the algorithm performance we have also explored the use of additional information that may be obtained by the radar on a sporadic basis. These sporadic signatures include SAR imagery, which may be obtained if a target stops while under track (a situation notorious for breaking moving-target indicator [MTI]–based tracks with a typical radar minimum detectable velocity of a couple of meters per second), and inverse SAR (ISAR) imagery, which may be obtained if a target turns while under track. Combined with the readily available HRR signature, these two-dimensional target signatures enable the target classifier and feature-aided tracker to perform robustly under a wide variety of circumstances.

The present work on identification and feature-aided tracking builds on significant previous developments at Lincoln Laboratory. The use of HRR signatures for target classification, which forms the cornerstone of moving-target classification, was originally demonstrated by M.L. Mirkin [1]. His results showed that HRR automatic target recognition (ATR) performance benefits from a number of independent looks at the targets, as well as fine resolution and extended dwell time (i.e., multiple looks for speckle reduction at each observation aspect angle). This work pointed out the reliance of HRR-based classification on effective tracking, since multiple looks on the same target will accumulate evidence only as long as all reports correctly associate with the track.

While Mirkin's work focused on HRR profiles for moving targets, L.M. Novak et al. performed classification using SAR imagery [2]. Their work on the Defense Advanced Research Project Agency (DARPA)–

sponsored Semi-Automated Image Intelligence Processing (SAIP) program demonstrated the successful application of a superresolution method called high-definition vector imaging (HDVI) [3], developed by G.R. Benitz, to SAR images. Their results indicated that a two-dimensional ATR system using HDVI-processed SAR images will outperform those systems whose SAR images were processed with standard non-data-adaptive processing methods (i.e., the weighted fast Fourier transform, or FFT).

Successful application of template-based ATR methods to ISAR imagery was later demonstrated by R.L. Levin [4], who used SAR images as templates. His results, in addition to demonstrating that ISAR images could be successfully used for moving-target ATR, also showed the benefits of the two-dimensional ISAR signature over the one-dimensional HRR signature, as well as the benefit of accumulating classification information over several looks, as a target executes a turn.

On the basis of the success of HDVI-processed SAR images in two-dimensional ATR systems, D.H. Nguyen et al. extended the application of HDVI to HRR profiles and demonstrated that the modified one-dimensional HRR ATR system (based on the system described in Reference 1) using these HDVI-processed HRR profiles provided significantly improved target recognition performance compared to other ATR algorithms using HRR profiles obtained from conventional image processing techniques, e.g., weighted FFT [5–7]. However, these HRR profiles were formed from SAR images of targets taken from the high-quality Moving and Stationary Target Acquisition and Recognition (MSTAR) data set having a high signal-to-noise ratio (SNR) of approximately 35 dB. From a radar resource perspective, however, for an HRR ATR system to be operationally useful, the system must operate at a significantly lower SNR, in the range of 20 to 25 dB. Unfortunately, the demonstrated performance of the HRR ATR classifiers suffered significant losses as the SNR decreased from 35 to 20 dB. In particular, the improved performance obtained with HDVI over conventional processing seen at the high SNR was significantly reduced and essentially eliminated at the lower SNR.

Recently, Nguyen et al. demonstrated the applica-

tion of Benitz's new superresolution technique called beamspace high-definition imaging (BHDI) to HRR profiles and showed that the enhanced one-dimensional ATR system using BHDI-processed HRR profiles exhibits significantly improved target-recognition performance at a low SNR compared to the conventional FFT method and the HDVI HRR method [8]. With the HRR profiles operating at a reasonable SNR, the radar resource requirements are sufficiently modest for frequent HRR interleave with MTI for effective feature-aided tracking.

This article is organized as follows: we first describe how the range profile, which is the only distinctive radar signature for moving vehicles that does not depend on the vehicle performing certain maneuvers, can be used to identify the vehicle class. In this context we introduce the concepts of improving the signature by using adaptive signal processing techniques (HDVI and BHDI). We illustrate how to use the signatures in conjunction with a template library to obtain classification via a Bayesian classifier. Once identification is established, we show how to use the vehicle classification so obtained, along with the radar HRR signature, to improve on track/report association, the feature-aided tracking component proper. Finally, we show how to extend this framework to include distinctive two-dimensional signatures, including SAR (stopping) and ISAR (turning) target signatures. These two-dimensional signatures perform both the functions of improving classification, by comparison with a template library, and they can also be used for on-the-fly template training in the absence of a pre-stored template library.

Moving-Target Classification

The main goal of any ATR algorithm is to correctly identify an unknown target from its remotely sensed signature. Given a sensed signature from an unknown target, many ATR systems work by matching the given signature against a set of candidate target hypothetical possibilities that could give rise to this signature. The technique we have used for ATR is to compare a radar signature with a template set from known vehicles.

With a high-resolution radar, there are two signatures that are readily available from moving vehicles,

the high-resolution range profile, or HRR, and an inverse SAR image, or ISAR. The HRR signal is a one-dimensional measure of the radar cross section along the range dimension of a vehicle. The ISAR image takes advantage of the vehicle turning to create a two-dimensional plan view of the vehicle. Other returns from moving parts such as lug nuts have been explored, but they are not considered in the present work for ATR purposes. If a vehicle stops, then it is also possible to obtain a SAR plan view, which can also be used for identification.

Since moving vehicles most often travel in relatively straight lines, it is usually impractical to obtain any vehicle signatures from a radar other than simple HRR profiles. The general approach to classification is to compare an HRR signature with a pre-stored set from known classes, and estimate the likelihood that each measurement corresponds to a particular vehicle class by using a matching algorithm. The better the match, the more likely the vehicle corresponds to a particular class. As the radar views the vehicle in subsequent scans, the classification evidence is accumulated for greater confidence.

The HRR ATR process proceeds over several stages. The basic profile is most directly formed by illuminating a target with a high-resolution waveform, for a series of coherent pulses (a coherent processing interval, or CPI). The target signature typically can be contained in a single Doppler bin for relatively short pulse streams by Fourier-transforming the pulse stream, and correcting for target range rate. Depending on the radar parameters, clutter cancellation may be necessary to better separate the target signature from the stationary ground clutter. Subsequent CPIs, collected adjacent in time, and perhaps with a center frequency offset, can be noncoherently combined to average out the speckle that arises from phase interference of different scattering centers on the vehicle. The methodology used for forming a range profile from a series of CPIs involves nominally averaging the power returns for the Doppler bins containing the target. We demonstrate an improved methodology for this profile formation, using superresolution methods.

Once a range profile has been formed from a radar detection, the process of classification involves matching it against profiles taken from a template da-

tabase of known vehicle types. A weighted mean square error (MSE) metric is used to measure the differences between the input test profile and the template profiles. A search over unknown parameters, such as precise target aspect angle, location, and radar cross section, is performed, and the smallest MSE is chosen to represent the fit between the test profile and the class being considered. For a multiclass classifier, we measure the MSE between the test profile and template from each class under consideration. This set of measurements—the MSE match between a given profile and each classification target set—is referred to as a *feature vector*.

Given a feature vector for a particular profile, it is necessary to infer actual classification. First the classifier is trained with data from known vehicles to obtain statistics of MSE distributions for both in-class vehicles and out-of-class vehicles. These training statistics are used to generate a likelihood that each measurement would come from a given class. Starting from a fixed *a priori* probability, successive likelihoods for each vehicle are combined with a Bayesian classifier to generate a posterior probability vector, which is presented to an operator.

Radar Resources and System Consideration

Moving-target ATR and tracking is a resource-constrained problem, and we cannot enter into a discussion of the topic without first delineating the primary issues involved. A radar is an active sensor, with a finite power and aperture. For a surveillance mission particularly, we must be extremely parsimonious with the radar resources.

The key radar resources to be minimized are the SNR, the number of pulses dedicated to a particular task, and the resolution of these pulses. Dedicating more pulses leads to a higher net SNR, but at the expense of less rapid revisits. Operating at very high resolution may require higher SNR, and possibly require additional pulses to cover a wider bandwidth. Moving-target classification generally functions in conjunction with an MTI mode, which operates by using a coherent pulse stream yielding a net SNR on the order of 15 dB. If an ATR system were to require 35 dB, or even 25-dB SNR to function effectively, and required that level of SNR for the entire coverage

area, the ATR system would need so much of the radar energy as to drastically reduce the MTI coverage area or drastically reduce the revisit frequency.

Superresolution Methods

To improve the features of the HRR profiles, given limited radar resources, we consider superresolution approaches. Unlike conventional Fourier imaging, which uses a predetermined set of weighting coefficients for each pixel in the output image, superresolution techniques are data-adaptive image-formation methods that select an optimal set of weighting coefficients for each pixel.

Superresolution methods have been used extensively for SAR ATR applications. G.J. Owirka et al. [9] used a two-dimensional template-based ATR metric to compare SAR images formed from several spectral estimation techniques and found that, of these methods, Benitz's HDVI [3] provided the best performance. These methods, including HDVI, consist of D.H. Johnson's eigenvector-based method [10], H.C. Stankwitz et al.'s spatially variant apodization (SVA) method [11], Stankwitz and M.R. Kosek's superresolved SVA (SSVA) method [12], and J. Capon's maximum likelihood method (MLM) [13].

Recently, Benitz proposed a modification to HDVI, which would reduce the computational costs from 8000 operations per input pixel for a SAR image by a factor of six. This new method, called beamspace high-definition imaging (BHDI) [14], provides speckle noise suppression that is comparable to its HDVI predecessor but improves on HDVI's ability to preserve target edges. Figure 1 compares a BHDI-processed SAR image of a Scud transporter erector launcher (TEL) against the image of the same Scud TEL formed by using a Taylor-weighted FFT (baseline FFT). The inherent ability of BHDI to control the main lobe and the sidelobes results in significant reduction in speckle noise compared to the weighted FFT.

BHDI, as applied to HRR profile estimation, involves an estimation of the range profile, given a set of noncoherent complex range cuts. For additional detail on this estimation process see the appendix entitled "Beamspace High-Definition Imaging Superresolution Method." Unlike the original SAR applica-

tion, the covariance estimation is aided by independent profiles from adjacent CPIs, so that the rank deficiency in the covariance matrix, which necessitates extensive constraints in the SAR estimation problem, is greatly alleviated.

Mean-Square-Error Matching Metric

The principles behind template-based ATR involve measuring the similarity between a test profile and what is expected from a given hypothesized target. The degree of similarity is used to determine the likelihood that the test profile originated from the hypothesized target class. For the ATR effort described here, we have employed a weighted MSE metric between the profile under test and the template, or expected profile, for a given hypothesized target.

Constructing a weighted MSE involves assuming a mathematical model of the profile under test. Since each range sample on an HRR range profile comprises a complex sum over all scatterers across the target with the same range, a single-CPI complex range profile is generally modeled as a zero-mean Gaussian random variate, with variance equal to the mean cross section as found in the template. Noncoherently combining several profiles results in a χ^2 variate. With such a model, an appropriate matching metric for a

profile to a template is an MSE metric between the profile and template, both expressed in decibels (dB). Several extra variables must also be included in the match, namely, matching of the receiver noise floor, a search over target range and radar cross section (the latter assuming that radar-cross-section calibration is not well known), and finally a search over template aspect angle to cover uncertainties in aspect-angle estimation. Figure 2 shows the effects of matching radar cross section. For an N -class classifier, we construct the MSE match between each test profile and all target classes being considered, which results in an MSE vector of length N .

Bayesian Classifier

With the MSE as the basic metric for similarity between each profile under test and each hypothesized target class, the task of classification involves an inference step from a measured MSE vector; i.e., given an MSE vector, how likely is it that the target under test corresponds to one of the target classes in the template set? We have employed a form of Bayesian inference, which constructs a posterior probability from a combination of prior expectations and conditional measurements, or likelihoods.

Successful application of Bayesian inference in-




FIGURE 1: Comparison of a beamspace high-definition imaging (BHDI)-processed synthetic aperture radar (SAR) image of a Scud transporter erector launcher (TEL) against a baseline image of the same Scud TEL formed by using a conventional two-dimensional Taylor-weighted fast Fourier transform (FFT). The Taylor-weighted FFT image on the left has 0.5 m $\times$ 0.5 m resolution. The image on the right has been BHDI-processed to an approximate 0.25 m $\times$ 0.25 m resolution. The BHDI-processed image is much sharper than the FFT-processed image as a result of speckle noise reduction.




FIGURE 2. The graph on the left shows the resulting test and template profiles after they have been correctly aligned in range and coarsely aligned in radar cross section (RCS). The graph on the right shows the resulting test and template profiles after completion of the RCS alignment procedure, with a correspondingly lower value of the mean square error (MSE) metric.

volves knowing something about the statistics of the MSE measurements. What is needed is the probability that a given MSE measurement is produced from a known target class. To construct such a probability, we employ training statistics, passing profiles from known vehicles through the MSE calculation and compiling statistics. Since the MSE vector is positive definite, and not easily modeled, we have computed statistics on the log of the MSE, which enables the MSE vector to be modeled statistically as a multivari-

ate Gaussian distribution. Figure 3 illustrates the process of forming log-MSE statistics and constructing a multivariate Gaussian model. In this figure we compile log-MSE statistics for M1 and M109 test target vehicles and a set of other vehicles, using templates for the M1 and M109.

The training statistics themselves represent the results of computing log-MSE values for a number of vehicles, over a number of conditions chosen to represent a test ensemble. Each point in the figure is the




FIGURE 3. Training statistics for a two-class classifier, in which the template set consists of M1 and M109 test target vehicle templates. The *other* class in the figure on the left consists of distinct vehicles similar in size to the test vehicles. The probability distribution function was modeled as a multivariate Gaussian distribution of the log of the MSE.

log-MSE for a given vehicle as measured against each of the filters for the test set. Figure 3 represents a two-class classifier, so each measurement is a point in the two-dimensional log-MSE space. With the ensemble suitably chosen, we arrive at a cloud of points representing MSE measurements under similar conditions.

These results illustrate a number of key points. First, even a vehicle matched to the template set yields significant scatter in log-MSE. This scatter is due in part to the intentional use of distinct vehicles for template construction and training, and in part to the intentional mismatch in collection geometry. These intentional mismatches are intended to represent the variability that we might encounter in practice when trying to match a profile collected in the field with previously stored templates. Another key point is that, in order to apply Bayes rule, we must estimate a probability distribution function for the scatter of measurements. Since the MSE is positive definite, it is not easily modeled by a simple estimation procedure. Modeling instead the log of the MSE alleviates this problem, and for purposes of classification a multivariate Gaussian appears to be a reasonable representation of the MSE vector for known vehicles.

Decision Regions

Figure 3 suggests that the training statistics for a given target class in an N -class classifier can be readily represented by an N -dimensional Gaussian random variate. The training statistics lead to a functional form for each target class, as represented in Figure 3 by the contours for the estimated Gaussian forms. Now, when an unknown profile is encountered, we first calculate a log-MSE vector for the profile against each of the test templates. Then, by using the training statistics, we evaluate the likelihood that the profile originated from each of the test target types, as well as the likelihood that the test profile corresponds to a vehicle outside the classification set. Application of the well-known Bayes rule [7] then converts these likelihoods, using prior probabilities, into posterior probabilities that the test template originated from each of the classification types or from an out-of-class vehicle.

With Bayes rule, the prior probabilities can strongly determine the estimated posterior probabilities. First, let us continue with the M1-M109 two-

class classifier as an example. There are two filters—M1 and M109—but there are three questions that we need to ask: is the target under test an M1, an M109, or neither? Setting the prior probability of neither (designated the *other* class) to zero forces the classifier to choose either an M1 or an M109 for any target profile under test. With equal prior probabilities for each of these two classes, Figure 4 shows the posterior probability that the target is an M109, as a function of the log-MSE measured from a given profile. The separation between the target decision regions is clearly evident.

Generally, we do not want to force a decision between one or another among a small number of target classes. More likely, most vehicles under observation are not one of the types being classified. By varying the prior probability of the other class, we can reject such vehicles. Figure 5 shows the effect of varying the *other* prior probability, or *prior*, on the decision regions. For even small *other* priors, MSE vectors that lie outside the main distribution of the M1 and M109 training distribution are rejected, allowing the classifier to correctly reject targets outside of the vehicle classes for which it has templates.




FIGURE 4. Posterior probability for the M1, assuming that the *other* prior is zero. Setting the *other* prior to zero reduces the three-class decision regions into the decision regions for the M1 versus the M109. If a feature vector lies in the red region, we classify the test profile resulting in that feature vector as an M1. Similarly, if a feature vector lies in the blue region, we classify the test profile resulting in that feature vector as an M109.




FIGURE 5. The contour lines in this figure represent the decision boundaries between the three classes as *a priori* knowledge about the *other* class varies from 0 to 1. The center line, which corresponds to an *other* prior of 0, is the same decision boundary shown in the previous figure. As the *other* prior increases from 0 to 1, the decision region for the *other* class increases at the expense of the M1 and M109 decision regions. Note that the M109 decision region compresses much more quickly than that of the M1, indicating a higher risk for an M109 high range-resolution profile than for an M1 to be classified as an *other* type.

HRR ATR Measure of Performance

In this section, we present target-identification performance results of the ATR classification methods described in the previous section. To graphically demonstrate the performance of the ATR classifier, we borrow the concept of a receiver operating characteristics (ROC) curve from communication theory, as shown in Figure 6. The *y*-axis on the ROC curve displays the probability of correct classification (*Pcc*) of a target of interest, while the *x*-axis shows the probability of false classification (*Pfc*) of a target from the *other* class as a target of interest. The *Pcc* is the probability that the targets of interest are correctly classified among themselves. The *Pfc* is the probability of falsely classifying one of the targets in the *other* class as one of the targets of interest. Note that an upward shift of the ROC curve corresponds to an improvement in *Pcc*, while a shift of the ROC curve toward

the left corresponds to a reduction in *Pfc*. In other words, any shift toward the upper left corner pertains to an improvement in ATR performance.

Description of the Data

The raw HRR profiles used in these studies were formed from SAR imagery provided to Lincoln Laboratory by Wright Laboratories, Wright-Patterson Air Force Base, Dayton, Ohio. These data were collected in 1995 and 1996 by the Sandia National Laboratory X-band (9.6 GHz) HH-polarization SAR sensor in support of the DARPA-sponsored MSTAR program.

The data set contains SAR images of forty-four military vehicles that were imaged in spotlight mode at 15° and 17° depression angles over 360° of aspect angles. Figure 7 shows ground-truth photographs of the twenty-three distinct MSTAR targets from collections 1 and 2 combined. The target set includes six distinct vehicles that are of interest to the classifier and seventeen vehicles that are not. We refer to the set of targets not of interest to us as the *other* class. Of the




FIGURE 6. Receiver operating characteristics (ROC) curves such as these are used to graphically display automatic target recognition (ATR) performance results. The probability of false classification (*Pfc*) represents the probability of falsely classifying targets not of interest as a target of interest. The probability of correct classification (*Pcc*) denotes the probability of correctly classifying among the targets of interest. Since the objective of the ATR is to improve the *Pcc* while at the same time reduce the *Pfc*, ATR performance results are improved by moving the ROC toward the upper left corner.

twenty-three targets in both classes, we have three variants each of the BMP2 armored personnel carrier and the M2 Bradley fighting vehicle, eleven variants of the T72 Russian tank, five variants of the M109 self-propelled howitzer, and four variants of the BTR70 armored personnel carrier. The variants of the M2, BMP2, M109, and BTR70 vehicles have minor differences between their configurations. The eleven T72 vehicles, however, vary significantly in configurations: (1) four T72s are fully intact without any fuel barrel mounted; (2) two T72s are fully intact with barrels mounted on the rear; (3) two T72s are missing skirts along the side but have fuel barrels mounted; (4) two T72s are missing skirts and have no fuel barrel mounted; and (5) one fully intact T72 with fuel barrels has reactive armor mounted on the exterior.

We used 17° depression-angle images of the six targets of interest taken from collection 1 to construct the six HRR classifier template profiles: BMP2#1, M109#1, T72#1, M1, M110, and M113. A total of

128 images of each of the six targets, covering 360° of aspect, were used to construct the HRR templates. For each SAR chip imaged over a 3° aspect look, we formed three 1° HRR templates by dividing the SAR image into three contiguous non-overlapping subsets. Each subset contains complex HRR profiles collected over a continuous 1° aspect look.

The HRR test profiles were generated from MSTAR target images taken at a depression angle of 15°. We tested the HRR ATR classifier on all forty-four targets, covering 360° of aspect angles. Targets taken from the *other* class were used as “confuser” vehicles and thus were not part of the template database. Ideally, we expect the ATR classifier to classify these confuser vehicles as an *other* type.

Summary of HRR ATR Results

We now summarize the effects of the SNR, the number of CPIs, and the range resolution on ATR performance results for HRR profiles conventionally pro-




FIGURE 7. The Moving and Stationary Target Acquisition and Recognition (MSTAR) collection 1 and 2 data set consisting of twenty-three distinct vehicles, six from the classifier target class and seventeen from the *other* class. Our six-class classifier is constructed from the M1, BMP2, M109, M110, T72, and M113 target vehicles shown in the upper left of the figure. This target set consists of three distinct models of the BMP2, four of the M109, and eleven of the T72. Of the eleven T72 vehicles, some have fuel barrels mounted on the back with skirts on the side, some have fuel barrels and no skirt, some have skirts and no fuel barrel, some have no skirt and no fuel barrel, and one has reactive armor mounted around the exterior. Four of these T72 variants are shown above. The *other* class of vehicles, representing targets we do not care about, consists of the seventeen targets shown on the right.

cessed by using a Taylor-weighted FFT, HDVI, and BHDI. Henceforth, we refer to the Taylor-weighted FFT as the baseline FFT method and HDVI as the baseline HDVI method.

First we show the gain in target identification performance with the new superresolution BHDI method, as compared against the conventional FFT-based method. Figure 8 compares the ATR results obtained with the BHDI-processed 20-dB SNR data at 1-m and 0.5-m range resolution against results obtained by using HDVI-processed and baseline FFT processed data. At 20-dB SNR the baseline HDVI and FFT methods provide comparable performance results. Figure 8(a), however, shows that the BHDI-processed data provides, on average, a 10% improvement in P_{cc} at all values of P_{fc} compared to data processed with the baseline HDVI and with the baseline FFT. Figure 8(b) expresses a similar improvement in performance with the BHDI method at the 0.5-m range resolution. In the 0.5-m resolution case, we see that the BHDI-processed data provide approximately 5% improvement in P_{cc} over the baseline HDVI-processed data and 8% improvement over the baseline FFT-processed data.

Moreover, at 20-dB SNR, the BHDI-processed data at 1-m range resolution exhibits performance

comparable to the baseline HDVI-processed and baseline FFT-processed 0.5-m range resolution data. In light of these results, we can expect a performance improvement equivalent to doubling the resolution by applying BHDI, compared to the two baseline methods.

In addition to providing improved performance at a lower SNR, BHDI also reduces the degree of system resources needed, as measured by, for example, the number of CPIs needed to form an HRR profile. Figure 9 compares BHDI-processed profiles with baseline FFT-processed profiles at 20-dB SNR, using five and three CPIs. The results indicate that when the number of CPIs is reduced from five to three, the probability of correct classification with the baseline FFT-processed HRR profiles degraded by 8% on average. With BHDI-processed HRR profiles formed from three CPIs, we obtained a performance result equivalent to the baseline FFT-processed profiles formed from five CPIs. Effectively, when BHDI is used, we can expect a performance improvement equivalent to nearly twice the number of CPIs used to form the HRR profiles, as compared to the baseline FFT method.

BHDI also saves radar system resources compared with HDVI. Initially, BHDI was developed to save




FIGURE 8. Classifier performance results for (a) BHDI-processed 1-m and (b) 0.5-m HRR profiles at 20-dB SNR processing compared against data processed with baseline HDVI and baseline FFT. The profiles were obtained by processing five coherent processing intervals (CPI). The BHDI-processed data provide approximately 5% improvement in P_{cc} over the baseline HDVI-processed data and 8% improvement in P_{cc} over the baseline FFT-processed data at 0.5-m range resolution.




FIGURE 9. Classifier performance results for (a) BHDl-processed 1-m and (b) 0.5-m HRR profiles formed from three and five CPIs, compared against data processed with the baseline FFT method. The test profiles were obtained at 20-dB SNR.

on computational requirements, i.e., reduce the number of operations needed to form an image. In addition, however, we have found that BHDl reduces radar resources while providing improved target-identification performance. For example, Figure 10 shows the performance results comparing BHDl-processed 20-dB and 25-dB SNR data at 1-m and 0.5-m range resolution against data processed with the baseline HDVI method. The results show that at 25-

dB SNR, BHDl provides approximately 5% improvement in Pcc at all values of Pfc. Moreover, at both the 1-m and 0.5-m range resolution, the application of BHDl at 20-dB SNR provides a performance equivalent to the data at 25-dB SNR processed with the baseline HDVI method. Consequently, we can obtain a performance improvement equivalent to 5 dB of SNR by applying BHDl, compared to the baseline methods.




FIGURE 10. Classifier performance results for (a) BHDl-processed 1-m and (b) 0.5-m HRR profiles at 25-dB SNR and 20-dB SNR compared against data processed with the baseline HDVI method. The test profiles were obtained from five CPIs.

Figures 8 through 10 contain results for six classes of vehicle types whose lengths and widths vary from 6.5 to 7.5 m and 3 to 4.5 m, respectively. To identify these targets among each other by using HRR profiles with such comparable range extents, the one-dimensional classifier must rely solely on the scattering features from the profiles. In the absence of confuser vehicles (*other* targets), we can expect the Pcc to reach 0.8 and 0.9 for the 1-m and 0.5-m range-resolution profiles at 20-dB SNR, respectively. The results demonstrated in Figures 8 through 10 include the presence of confuser vehicles with lengths and widths comparable to the six classes under test. Unfortunately, identifying targets with such comparable range extent with high confidence is difficult. We therefore expect the Pcc to deteriorate significantly as the Pfc of the *other* class improves.

We can experimentally verify this assertion by including, into the template database, HRR profiles of a Scud TEL and M978 Hemmt truck whose lengths and widths exceed 10 and 2.5 m, respectively. These targets have twice the range extents of the six vehicle classes used to generate the results shown in Figures 8 through 10. After testing the HRR classifier with the addition of the Scud TEL and M978, we can see in Figure 11 that the Pcc of the resulting eight-class classifier improves on the six-class classifier discussed earlier by approximately 10% in the regime of low Pfc (< 0.2). The improvement mentioned comes from averaging individual Pcc results from the original six classes with the high Pcc results from the Scud TEL and M978. The black and green curves show individual results for the Scud TEL and M978, respectively, indicating discrimination of these long targets well exceeding 95% Pcc. In the $Pfc < 0.01$ regime, the classifier performance degrades significantly as a result of setting the *other* prior probability close to one (hence targets' prior probabilities close to zero).

In many battlefield situations, only very high-value targets such as a Scud TEL may be of interest to the commander. Given this presumption, we may use the one-class classifier to discriminate a Scud TEL from the rest of the world. Figure 11 contains performance results of a Scud TEL discriminator using HRR profiles at 1-m range resolution and 20-dB SNR. At Pfc values greater than 0.3, the classifier can identify a

Scud TEL from the other targets almost perfectly. However, the performance degrades considerably as the Pfc improves to below 0.3. We suspect two contributing factors to the observed degradation. First, the M978 Hemmt acts as a confuser vehicle, very similar in range extent to the Scud TEL. With nothing but training statistics for different vehicles, the classifier does not have a good criterion for rejecting the M978. Second, near broadside the MSE values from all vehicles tend to collapse into a common value, so that classifier performance is much poorer for near-broadside aspects than for other aspects.

Other Class Information

At this point, we need to make a slight diversion back into Bayes decision theory. We have come across the situation in which a confuser vehicle prevents the classifier from simultaneously classifying a Scud TEL as a TEL, and rejecting a similar length M978 as a non-TEL. The solution is twofold. The first part of the solution is, not surprisingly, that it is necessary to




FIGURE 11. Classifier performance results for BHDl-processed 1-m HRR profiles at 20-dB SNR. The red curve represents previous results obtained for the six-class classifier (with targets M1, T72, M109, M110, BMP2, and M113), while the blue curve shows the eight-class classifier performance resulting from inclusion of the Scud TEL and the M978 into the six-class classifier. The black and green curves show individual results for the Scud TEL and the M978, respectively, indicating that the classifier has no trouble discriminating these two new targets from vehicles in the *other* class.

add a template for the confuser, if it exists, to the classifier set. If a precise template is not available, it is still beneficial to generate a surrogate confuser to use in the template set. The second part of the solution will involve modifying how the Bayes decision rule is handled.

Having a confuser template adds information that we would expect to improve the ability to distinguish between a target of interest and a target that one wishes to reject from classification. This additional information plays a role in the Bayes decision process by adding additional dimensionality to the problem. Rather than simply rejecting a target as an *other* class, which is rejection based only on targets we wish to classify, we can reject a target on the basis of the additional information that the target under test in fact yields a good MSE match to a confuser template. When handled correctly, such information always improves the ability of the classifier to perform.

In evaluating the efficacy of adding additional information to the classification, we need to be careful to keep the criteria for success and failure constant as we add information. This careful attention plays in two places in the decision process—when formulating the decision and when scoring the result. The decision process is first obviously changed by the addition of the other class template, as we would expect. The decision process is next changed by splitting the *a priori* probability for the other class between the targets without templates and the confuser for which a template is provided. The second place in which the additional information affects the decision process is more subtle but, in hindsight, more obvious: the criteria for success should not change. Thus, for example, if we consider a single-class Scud TEL detector and add an M978 filter to aid in the rejection of this similar length vehicle, then the criterion for success will still be whether the TEL is properly classified as a TEL, and if a non-TEL is properly identified as a non-TEL. If an M978 is improperly classified as an *other* type of vehicle, or an *other* type of vehicle is improperly identified as an M978, then this classification is not an error.

When the classifier is provided additional information about the M978 by using the approach described above, we see that the classifier Pcc performance im-

proves dramatically, as shown in Figure 12. For a simple single-class classifier, for the example shown, the classifier achieves a probability of detection of 80% while correctly rejecting 80% on non-TELs. Here, note that for a single-class classifier, the probability of detection is synonymous with Pcc. With inclusion of the M978 filter, the 80% probability of detection is achieved while rejecting 90% of non-TELs. As more and more confuser filters are included, the classifier approaches perfection. Additionally, since the widths of all the vehicles are comparable, especially on the scale of the 1-m resolution used, significant and unavoidable confusion occurs at broadside looks. When classification in the context of a tracker is applied, the classifier does not effectively update near-broadside looks when all targets are indistinguishable. Simulating this situation for the example here, we exclude targets near broadside, and find the simple TEL detector can achieve a 90% probability of detection with a 93% rejection of *other* vehicles, and adding the M978 improves this result to 90% probability of detection and 97% rejection of non-Tels.




FIGURE 12. The dotted blue curve shows the results of the one-class classifier trying to discriminate a Scud TEL from all other targets. When provided with additional knowledge of the M978 vehicle, the classifier performance improves by 20% in the Pfc < 0.2 regime. If profiles within $\pm 30^\circ$ of broadside are removed, the performance improves by approximately 35% in the Pfc < 0.2 regime. As knowledge of additional vehicle targets are added to the classifier, the Scud TEL discriminator performance approaches the desired performance result.

Feature-Aided Tracking

For moving targets, target identification cannot be considered separately from tracking, since the targets are in motion, and target identification becomes rapidly perishable information if the target is not kept in track. As was discussed above, there are three primary target signatures that are useful for identification: (1) the HRR range profile that is visible when the target is in motion, (2) the ISAR signature that is visible when the target undergoes a turn, and (3) the SAR signature that is visible when the target stops. For identification using HRR, several looks at the target from different aspect angles are necessary to arrive at a high-confidence identification, requiring the tracker to maintain target identity over the accumulation period for the classification evidence to be properly accumulated. Since ISAR and SAR both yield a two-dimensional signature, we would expect a greater classification confidence than with HRR for each observation. But again the tracker must properly link reports and maintain track for such identification to be of anything but a transient nature.

We have pursued the viewpoint that identification information can be accumulated from a number of signatures, be it HRR, ISAR, or SAR. Linking this information together requires a tracker that not only can fuse different identification types such as HRR, ISAR, and SAR, but ideally one that can also take advantage of this non-kinematic information to improve tracking.

The objective of a tracker is to maintain track on moving targets, and to estimate the target state. A kinematic tracker produces estimates of the track state (e.g., position, velocity, and heading) based on radar measurements such as range, range rate, angle, and angle rate, by correctly associating new detections with existing target tracks assuming appropriate kinematic transition models. A feature-aided tracker builds upon the traditional kinematic tracker and uses additional target features, such as the HRR profile and target classification information, to help improve this association.

Feature-aided tracking encapsulates a suite of algorithms, consisting of signature-aided tracking (SAT), classification-aided tracking (CAT), and classification

diffusion-aided tracking (CDAT). SAT uses the degree of similarity between the HRR profiles obtained from successive detections of the target. No inference is made in regard to the identification, or type, of the vehicle under track. In other words, SAT addresses the question of whether the target at the current time is the same as the target from an earlier time. In contrast, CAT uses the degree of similarity between the identification associated with the current report and an existing track; i.e., if the identification of the report and track are similar, then the report and track are more likely to be associated. Finally, CDAT attempts to mitigate the degradation in target identification that occurs when the tracker associates a report with the wrong track [15]. The approach of CDAT is to blend prior classifications from tracks that are kinematically ambiguous so that built-up classification information is not completely associated with the wrong report.

This article focuses on the more mature feature-aided tracking methods, namely, SAT and CAT, and defers further discussion of CDAT until a later date. In the next four sections we describe SAT and CAT in general terms and present initial results that strongly suggest that sufficient additional benefit can be obtained from feature-aided tracking.

Signature-Aided Tracking Algorithm and Architecture

For the purposes of feature-aided tracking analysis, we use a kinematic tracker previously developed at Lincoln Laboratory in the 1990s to support a small, lightweight radar payload located on an unmanned air vehicle, or UAV [16]. The focus of this article is not on the details of this particular tracker, but rather on how the feature-aided tracking algorithms fit within this very general type of tracker architecture. In general, a tracker receives M reports from a sensor and attempts to find the best match between these reports and N existing target tracks. (For simplicity, we ignore any discussion of how the tracks are initiated and begin our discussion with the tracks already in existence.) We briefly walk through the basic steps of a generic kinematic tracker to make clear how the feature-aided tracking algorithms are added.

First, for a given hypothesized track-report pair, the tracker determines if the track and report are suf-




FIGURE 13. Tracking of two targets—an M1 (upper left) and an M109 (lower left), where each starts out moving on a separate road. The goal of the tracker is to associate the M1 track to the M1 report and the M109 track to the M109 report. This association problem is unsolvable, given only the targets' kinematics information. We can supplement the kinematics information with the targets' HRR profiles to provide target identification and resolve the ambiguity in the kinematics.

ficiently “close” in a kinematic sense such that they may be plausibly associated. If so, the track is propagated (i.e., extrapolated) to the time of the report. A non-negative score is then calculated on the basis of the degree of similarity between the kinematic states of the track and report. This process is repeated for all tracks and reports from the current radar scan, generating a matrix of scores—one number for each plausible track-report pairing. To determine the correct associations, we use an optimization algorithm to find the best pairings among all the viable hypotheses. We now discuss the problem that SAT tries to solve, our approach for solving this problem, and the architecture for implementing this solution.

Consider the simple scenario shown in Figure 13 as an illustration of the problem, where, at time t , as shown at the left, two targets are widely separated and located on different branches of the road. We arbitrarily designate one of the vehicles as an M1 (shown in blue), and it is located on the top branch of the road with aspect angle θ_1 . The other vehicle we design-

nate as an M109 (shown in red) with aspect angle θ_2 . The corresponding HRR profiles for the two vehicles are also shown, with the scattering differences between the two clearly visible. At a later time, the two vehicles have merged on the road. Using kinematic information alone, we have an approximately 50% chance of obtaining the correct association. However, if we use the HRR profiles obtained from the targets at time $t + \Delta t$, as shown on the right in Figure 13, and compare with the target profiles at time t , we expect a greater chance of getting the association correct; i.e., the degree of similarity between the track and report HRR profiles will help resolve the ambiguity of the track-report association. The degree of similarity between the two profiles at $t + \Delta t$ with the two profiles at earlier time t (four possible cases) is used to help determine whether the M1, for example, is the lead or trailing vehicle at time $t + \Delta t$. We now discuss the particulars of the algorithm in some detail and explain how to modify the kinematic tracker for SAT.

The main assumption of SAT is that if the HRR

profiles of the track and report indeed come from the same vehicle, then the corresponding differences in the profiles, as measured by the MSE score, should be smaller than if the HRR profiles came from two different vehicles. How do we characterize these differences? Using the DARPA MSTAR vehicles, we computed MSE scores for all possible pairings of vehicles for all aspect angle differences less than 5° . (We found that for angles greater than approximately 5° it is more difficult to distinguish between vehicles with similar sizes only on the basis of their HRR profiles). The HRR profiles used to generate the above mentioned results have 1-m range resolution and 20-dB SNR, and were formed by using the superresolution BHDI method developed by Nguyen et al. [8] at Lincoln Laboratory.

We observed that the so-called *matched* scores (the vehicles are the same) and *mismatched* scores (the vehicles are different) are each normally distributed with the match mean, as expected, smaller than the mismatched mean. The variance of the mismatch case was slightly larger than the match value. In the following discussion we represent the match and mismatch mean by μ_{match} and $\mu_{mismatch}$, respectively, and the match and mismatch standard deviation by σ_{match} and $\sigma_{mismatch}$, respectively. Consequently, given an

MSE score for a track-report pair, we can calculate the log-likelihood ratio by using the expressions for the probability density function (pdf)

$$pdf_{match}(x) = \frac{1}{\sqrt{2\pi}\sigma_{match}} e^{-\frac{1}{2\sigma_{match}^2}(x-\mu_{match})^2}$$

$$pdf_{mismatch}(x) = \frac{1}{\sqrt{2\pi}\sigma_{mismatch}} e^{-\frac{1}{2\sigma_{mismatch}^2}(x-\mu_{mismatch})^2}$$

to obtain a signature χ^2 score,

$$\chi_{signature}^2(x) = -\frac{1}{2} \left\{ \frac{(x - \mu_{mismatch})^2}{\sigma_{mismatch}^2} - \frac{(x - \mu_{match})^2}{\sigma_{match}^2} \right\}.$$

The signature χ^2 value is then added to the kinematic-based χ^2 score to obtain the total SAT χ^2 score.

Figure 14 shows the Gaussian probability density functions and corresponding χ^2 scores as a function of MSE scores. As expected, the signature χ^2 score increases monotonically with MSE score. Recall that the larger the χ^2 score, the less likely the track and report signatures belong to the same target.




FIGURE 14. The left graph shows the match and mismatch Gaussian probability density functions obtained from training HRR profiles at 1-m range resolution and 20-dB SNR. We generate the training data used to estimate the match probability density function by comparing profiles from the same target but separated by up to 5° in aspect. To generate the training data for the mismatch probability density function, we randomly compared profiles from two different targets. The negative log-likelihood ratio of match to mismatch is used to normalize an MSE score into a signature χ^2 value. The right graph depicts the transformation of the log-MSE scores into equivalent χ^2 values, using the match and mismatch mean (μ) and standard deviation (σ) distributions shown in the left graph. As expected, the χ^2 values increase monotonically with log-MSE scores.




FIGURE 15. The key components of the signature-aided tracker with HRR profiles as features. For each hypothesized track-report association, a maneuver is detected and a kinematics χ^2 value computed from the extrapolated track state. Each track and report contains, in addition to the kinematics information, an HRR profile with the track also containing an estimated aspect angle of its HRR profile. The Kalman filter then updates the state of the track to the hypothesized report from which the report-HRR aspect angle is estimated. If the track and report-HRR profiles differ by a large aspect angle, then the feature information is ignored. Otherwise, a signature χ^2 is computed and fused with the kinematics χ^2 .

We now present the architecture for implementing the SAT algorithm. Figure 15 shows a high-level diagram of this architecture. As was discussed with the kinematic tracker, we assume that there are N existing tracks and M new detections (reports) at time t . Furthermore, for the sake of clarity and simplicity, we assume that each track and report has an HRR profile. Consider first a single hypothesized track-report pair. As before, using kinematic information alone, we determine if the track and report are sufficiently close so that the hypothesized pairing is plausible; i.e., are the range, range rate, and angle, for example, sufficiently similar? If so, we project the kinematic state of the track forward to the time of the report and we compute the kinematic-based χ^2 score. Since HRR profiles are quite sensitive to target aspect angle, resulting

in a loss of target specificity after an aspect-angle change of approximately 10° , we need to exclude those signature scores computed by matching track and report signatures differing by more than 10° in aspect. To properly determine the aspect angle of a report, the tracker must calculate an updated estimate of the track aspect angle. We therefore assume for the moment that this track does associate with the report and consequently we filter the report data with the track to obtain an updated estimate of the track state.

Using the updated state estimate, we determine the aspect angle of the report and ask the following question: has the aspect angle between the track and the report changed too much such that SAT should not be used? If the answer is yes, then we, as with the traditional kinematic tracker, use only the kinematic in-

formation to determine the score associated with this track-report pair. If, however, the track has undergone a change in aspect angle that is sufficiently small, then we calculate the signature χ^2 score, which is then added to the kinematic χ^2 value to obtain a total χ^2 score. Repeating this process for all possible track-report pairings, and using proper normalization of the scores to account for the different number of degrees of freedom, we obtain an association matrix from which the final track-report pairings are made by optimizing for this tracker with Munkres algorithm.

Discussion of SAT Performance Results

We now present some initial SAT results that indicate the level of improvement to expect in tracker performance. To obtain these results, we use instrumented ground vehicles, simulated target detections, and HRR profiles derived from SAR imagery. The vehicle motion data were collected for the DARPA Affordable Moving Surface Target Engagement (AMSTE) program at Patuxent River Naval Air Station, Maryland, in 1999, by using Global Positioning System (GPS)-instrumented vehicles. The sensor detections were derived from the GPS data by adding sensor noise and measurement error consistent with a typical radar sensor. The radar scan rate used to generate the detections was set to five seconds, which is consistent with a small-region-of-interest rapid-scan mode. The probability of detection was set to 0.9. The MTI detections were augmented with HRR profiles formed with the superresolution BHDI method [7] at 20-dB SNR with 1-m range resolution, using SAR images collected for the DARPA MSTAR program.

We consider a simple scenario consisting of two vehicles—an M1 and an M109—traveling along a straight road. For this scenario, there are 320 detections with 130 of these detections resulting in potential misassociations. Figure 16 illustrates one particular example of the type of confusion that may occur for this scenario. At times t and $t + \Delta t$ the M109 is in fact ahead of the M1. However, from the perspective of the tracker, an equally plausible alternate hypothesis exists: the M1 overtook the M109 during the time interval between scans and is now ahead of the M109. Figure 16 illustrates these two kinematically ambiguous hypotheses.




FIGURE 16. A sample two-target tracking scenario consisting of an M1 and an M109 following each other closely. Two position measurements of each vehicle are shown in this figure. Hypothesis 1 depicts the true trajectory of the targets, with the M1 following the M109. Alternatively, the M1 could have passed the M109 between the two position measurements, which would result in a reversal of track identification, as shown in hypothesis 2.

Using kinematic information alone, the tracker got the vehicle identification wrong 21% of the time. That is, 27 times out of 130, the tracker paired the reports with the wrong tracks. In contrast, by using SAT, the number of misassociations was reduced from 21% to 9%; i.e., the number of misassociations was reduced by more than a factor of two. Figure 17 illustrates these results. This simple example illustrates the potential for SAT to help the tracker resolve kinematically confusing close encounters between vehicles. Clearly, additional work—e.g., more challenging scenarios with an expanded set of tracker measures of performance, such as track lifetime and continuity—are needed before a more quantitative measure of the added value of SAT can be stated with full confidence. These results strongly suggest, however, that additional effort in this area is warranted.

One-Dimensional SAT Variants

We have seen how HRR profiles can be used to improve tracker performance. HRR profiles are, however, typically obtained from vehicles that are moving with a range rate above the minimum detectable velocity along a relatively straight line. When a vehicle turns or stops, it is possible to obtain two-dimensional signatures that are more informative than HRR




FIGURE 17. (a) The two-target simulation contains a total of approximately 450 detections made at five-second update intervals and a probability of detection of 0.9. Of the 450 detections, 130 of them could potentially result in misassociations. (b) By using only kinematics information, the kinematics tracker made 27 misassociations, or 20.8% of the 130 possible misassociations. In signature-aided tracker (SAT) mode, the number of misassociations dropped by over half to 12, or 9.2% of the possible misassociations.

profiles. In this section, we present two variants of SAT that make the best use of the vehicle signatures collected while the vehicle is turning or stopped. The proposed approach takes SAR and ISAR images, obtained while the target is stopped or turning, respectively, and collected at target aspect angle θ . We then generate HRR profiles at an angle $\theta + \Delta\theta$, where $\Delta\theta$ is large (i.e., much greater than the few degrees we would normally anticipate). The ability to generate HRR profiles at angles substantially different from the image collection angle results in a significant extension of the range of applicability of SAT.

We first consider the case of HRR profiles derived from a SAR image. The procedure for generating HRR profiles for angles present in the SAR image are

well understood and commonly used. What is new, however, is that by performing a simple rotation of the SAR image, we can obtain representative HRR profiles at angles outside the range of normally expected values. There apparently are target features (e.g., dominant scatterers) present in the two-dimensional image that are sufficiently robust to these types of operations (i.e., simple rotations), at least for the purposes we are considering (one-dimensional SAT). For example, in Figure 18, we show the distribution of match and mismatch MSE scores for HRR profiles against HRR profiles derived from SAR images. The match scores are formed from the HRR profile of a Scud TEL at angle θ against the HRR profile of a Scud TEL derived from a SAR image at angle $\theta + 30^\circ$. The mismatch scores are obtained by comparing the HRR profile of a Scud TEL at an angle θ against the SAR-derived HRR profile of a M60 tank at aspect angle $\theta + 30^\circ$. The HRR profiles are at 1-m range resolution and 35-dB SNR. For the case considered here, the separation between the match and the mis-




FIGURE 18. The match and mismatch probability density functions resulting from comparing HRR profiles of a Scud TEL against SAR-derived HRR profiles of a Scud TEL and an M60 tank, respectively. The aspect-angle separation between HRR profile and SAR image was approximately 30° . In the match case, SAR images of the Scud TEL are rotated to the aspect angle of the HRR profiles. By collapsing the rotated SAR image in range, we can form an HRR profile. The resulting HRR profile is then matched against the given HRR profile. In the mismatch case, we replace the SAR images of the Scud TEL with images of an M60. The HRR profiles used in this case have a range resolution of 1 m and SNR of 35 dB.




FIGURE 19. The expected match χ^2 values compared against the expected mismatch χ^2 values for aspect-angle differences from 10° to 180°. At all aspect-angle differences, the χ^2 probability density functions show a notable separation between match and mismatch.

match scores is clearly evident, and as a result it is expected that the SAR-derived HRR profiles will provide greater aspect-angle diversity, thus extending the applicability of SAT.

For example, as shown in Figure 19, the separation between the mean of the match and mismatch scores persists for a very large range of aspect angles, far greater than could possibly have been anticipated. Naturally, these results are most likely reflective of the great dissimilarity between the Scud TEL and the M60 tank, and we fully expect the separation to lessen as the vehicles become more similar in size, while the depression angle is reduced to values more consistent with long-range surveillance geometries.

The same approach can, in general terms, be applied to HRR profiles derived from ISAR images. Figures 20 and 21 show the corresponding results. Although this method is not apparently as robust as the SAR-based method, there is still considerable separation between the match and mismatch scores, which persists over the full range of angles from 0° to 180°.

In summary, target features in HRR profiles derived from SAR and ISAR images that have undergone significant rotations (to align them with the corresponding HRR profile) appear to be robust to angle diversity, thus significantly extending the applicability of SAT to more challenging scenarios.




FIGURE 20. The match and mismatch probability density functions resulting from comparing HRR profiles of a Scud TEL against ISAR images of a Scud TEL and an M60, respectively. The aspect separation between HRR profile and ISAR image was approximately 30°. We perform the match and mismatch comparison in the same manner as in the HRR versus SAR mode. However, the ISAR mode introduces an additional complication in that the cross-range resolution must be estimated prior to rotation. The MSE computation, in this case, requires a search in cross range for the appropriate resolution.




FIGURE 21. The expected match χ^2 values compared against the expected mismatch χ^2 values for aspect-angle differences from 10° to 180°. As shown, the mean match χ^2 increases with the aspect-angle difference. As the aspect-angle difference increases, the separation between match and mismatch χ^2 probability density functions decreases. For aspect-angle differences greater than 60°, however, the match distribution falls almost entirely within the mismatch distribution. Consequently, a χ^2 value near zero could have resulted from either a match or mismatch comparison.

Classification-Aided Tracking Algorithm and Architecture

Classification-aided tracking (CAT) uses information about the identification of the target to help the tracker resolve kinematic ambiguities [17, 18]. Unlike SAT, which uses the target signature directly to determine if the target at time $t + \Delta t$ is the same as the target at an earlier time t , CAT goes one step further and determines if the target at the two times is the same vehicle type. The degree in similarity between the identification of the track and the identification of the report is used to compute a CAT score, which, as with SAT, is added to the kinematics-based score. We can think of CAT as closing the feedback loop in the tracker with classification; classification information on targets is built up with successive reports associated and accumulated in the tracker. With CAT, the classification itself is used to improve the tracker association performance, which in turn improves classification. In this section we present the CAT algorithm and a high-level architectural diagram illustrating the implementation of the method, and we conclude with some preliminary but suggestive results.

As with SAT, CAT functions by modifying the pairwise χ^2 error between reports and tracks. The added term is referred to as the CAT score, with a small value designating a good match that increases the likelihood of association, and a large value indicating a poor match that decreases the likelihood of association. To calculate the CAT score we assume that we have an estimate of the identification of the target track; i.e., we have some prior knowledge of the track identification, built up from previous reports. The first step in calculating the CAT score is to compute an estimate of the identification of the target associated with the sensor report. However, to perform this calculation, we must first have an updated estimate of the track aspect angle. As we did with SAT, we obtain a hypothesized estimate of the report aspect angle by updating the track state with the report information, as if the track would associate with the report. We then match the report HRR profile with HRR profiles in the template library for angles near the estimated angle. We compute an MSE vector $\mathbf{M}$, where each element of the vector corresponds to an

MSE score between the report HRR profile and the best matching template HRR profile taken from each class type. By properly conditioning the MSE scores as described earlier, we can approximately fit a multivariate Gaussian distribution to the MSE vector whose first-order and second-order statistics have been pre-calculated during training. Using Bayes' rule, we can combine *a priori* knowledge of the track identification with the multivariate Gaussian likelihood generated from the MSE vector to calculate the *a posteriori* probability that the report corresponds to a target of a particular type. Given the set of *a posteriori* probabilities, referred to as the *classification vector*, we can now calculate the CAT score, as described below.

To determine if track i should associate with report j , we ask the following question: what is the likelihood that track i gave rise to (is connected with) report j on the basis of identification information? Basically, as the degree of similarity between the identification (as measured by the classification vector) of track i and report j increases, it is more likely that this particular pair should be associated. To compute this likelihood, we let the probability that track i gave rise to report j be given by the expression

$$P(\text{track}_i \rightarrow \text{report}_j) \propto \sum_{n=1}^{N+1} P_n(\text{track}_i) L_n(\mathbf{M} | \text{report}_j),$$

where $P_n(\text{track}_i)$ is the prior probability that track i is a target of type n , and where $n = 1, 2, \dots, N + 1$. (The $N + 1$ element is for targets of the *other* class.) The term $L_n(\mathbf{M} | \text{report}_j)$ is the likelihood that a target for the j th report gave rise to the MSE vector $\mathbf{M}$ when the j th report is compared with the templates from the N classes.

To illustrate the computation of the CAT score in simpler terms, we consider the following example. Let the degree of the classifier N be 1; i.e., we are interested only in determining if the target is of a single particular type, say a Scud TEL. Consequently, the $N + 1 = 2$ elements of the classification vector correspond to the probability that the target is a TEL or a non-TEL. For this case, the expression for the CAT score can be expressed as

$$P(\text{track}_i \rightarrow \text{report}_j)$$

$$\begin{aligned} &\propto (\text{probability that track } i \text{ is a TEL}) \\ &\times (\text{likelihood that a TEL gave rise to MSE score}) \\ &+ (\text{probability that track } i \text{ is other}) \\ &\times \left(\text{likelihood that an other gave rise to} \right. \\ &\quad \left. \text{MSE score} \right). \end{aligned}$$

Given the value of $P(\text{track}_i \rightarrow \text{report}_j)$, we compute the χ_{ID}^2 score as

$$\chi_{ID}^2 = -\log[P(\text{track}_i \rightarrow \text{report}_j)],$$

and subsequently add the term χ_{ID}^2 to the kinematic score $\chi_{kinematic}^2$ to obtain the CAT score χ_{CAT}^2 .

Figure 22 illustrates a high-level architectural diagram for the CAT algorithm. The steps are nearly identical to SAT in that a feature-based score is computed and added to the usual kinematic score. The difference from SAT, of course, is that the score is now derived from an estimate of the target identification. The steps of this block diagram have been described in detail in the earlier section entitled “Signature-Aided Tracking Algorithm and Architecture.”

Discussion of CAT Performance Results

We now discuss some preliminary results using CAT. We consider a six-class classifier, in which the vehicles of interest are M1, M109, T72, M110, M113, BMP2, and the *other* class, consisting conceptually of all vehicles that are not members of the six classes. As before, we use BHDI-formed HRR profiles at 1-m range resolution and 20-dB SNR. Likewise, we use the same radar simulation parameters (i.e., radar noise, measurement error, and a scan rate of five seconds) as was done with SAT. We first examine target identification performance and subsequently show how identification can be used to improve tracker performance, as measured by the number of correct track-report associations, with the CAT method.

For the purposes of assessing identification performance, we consider a single-target scenario and designate the target as an M109. The vehicle identification of each track is characterized by a classification vector. For example, the first element is the probability that the target is an M1, the second that the target is an M109, and so on. In the top graph of Figure 23, the accumulated *a posteriori* probability that the target is an M109 is shown as a function of scan number. A




FIGURE 22. The incorporation of the HRR ATR classifier into the kinematics tracker to form the classification-aided tracker (CAT). For each kinematically possible association between track and report, we compute the kinematics χ^2 score as described previously in Figure 15.




FIGURE 23. The top graph contains the accumulated posterior probability of classification for the M109 as a function of scan number. A blue dot in this graph indicates that the largest value in the classification vector for this scan corresponded to the M109. A red dot indicates that the largest value in the classification vector corresponded to the *other* class. On average, the probability of correctly classifying the M109 was 0.74. The second graph shows the aspect error resulting from estimating the M109's aspect angle from its heading. The third graph shows the estimated aspect angle of the M109. The bottom graph depicts the M109's true speed as a function of scan number.

blue dot indicates that the largest value of the classification vector is the element corresponding to the M109. In contrast, red dots indicate those times when the largest value of the classification vector belonged to the *other* class. For this case, 74% of the time the tracker correctly identified the target as being an M109.

To help understand the possible source of errors in classification, we examined the following parameters: the error in the aspect-angle estimate, the value of the aspect angle, and the speed of the vehicle. We anticipate that the identification performance will degrade as the error in the aspect-angle estimate increases. Likewise, as the target approaches broadside, we also expect performance to degrade because targets appear more similar when viewed from the side. Furthermore, as the speed of the vehicle approaches zero, the range rate will necessarily also approach zero, result-

ing in a poorer estimate of the target velocity and corresponding heading, ultimately resulting in a poor estimate of the target aspect angle. For the second and third graphs in Figure 23, we observe that when there is a large change in aspect angle, there is often a significant increase in the error of the estimated aspect angle, resulting in a decrease in target identification performance, as indicated by the increase in the number of red dots. For example, for scan numbers between 300 and 400, the vehicle is undergoing significant changes in aspect angle (third graph), resulting in large spikes in the error of the aspect-angle estimate (second graph), giving rise to a larger number of red dots (top graph). During this interval, the speed of the vehicle changed markedly, often nearing zero, which further exacerbated the estimation of the target aspect angle. And finally, although not clearly evident from this figure, due to the scale, it can be shown that




FIGURE 24. The accumulated posterior probability of classification for the M109, given perfect estimation of the report aspect angle, as a function of scan number. The blue dots indicate that the M109 was correctly classified with maximum *a posteriori* probability. The red dots indicate that the M109 was classified as an *other* vehicle. On average, the probability of correctly classifying the M109 was 0.83.

the identification performance did indeed degrade when the target was near broadside.

To estimate the expected improvement in tracker performance as the estimate in target aspect angle improves, we simply substituted the true value of the aspect angle for the estimated value; i.e., we assumed there was no error in the estimation process. This approach provides an upper bound on the gain in performance as related to errors in angle estimation. Figure 24 shows that there are fewer red dots indicating misclassification for this case, resulting in an increase in the percentage of time the tracker correctly identified the target as an M109 from 74% to 83%. Consequently, this change suggests that any improvements to the tracker that result in a decrease in the error of the aspect-angle estimate, such as added road constraints, can be expected to provide significant improvement in identification performance.

We now demonstrate how CAT, using target identification, can be used to improve tracker performance. We consider a scenario in which two vehicles approach an intersection from different roads. At the intersection, one vehicle (an M109) turns right and the other (an M1) turns left, as illustrated on the left in Figure 25. This results in a kinematically challenging situation. With the usual motion models, the illustration on the right in Figure 25 is a more likely hypothesis for this scenario, where each vehicle is assumed to continue along a straight path. Using kinematic information alone, the tracker chose the correct path 20% of the time; i.e., it determined that the vehicles had turned. In contrast, with CAT, the percentage increased to 100%—a significant improvement in tracker performance.

The CAT performance results, shown in Figure

26, show tracker improvement and suggests a strong performance dependency on *a priori* knowledge of each track's vehicle identification. In particular, as the prior knowledge of the identity of the vehicle increases, the separation in the values of target identification χ^2 corresponding to the correct and incorrect association increases, resulting in an increase in the probability of correctly identifying the target. For example, as the value of the prior increases from 0.5 to 0.9, the penalty for an incorrect association increases from -5 to -3 , while the benefit for a correct association increases from -7 to -8 (a lower score is better).




FIGURE 25. Consider a tracking scenario in which two targets—an M109 and an M1—moving on separate roads come together briefly at an intersection and then each turns and diverges. The left image shows the true motion of the vehicles (hypothesis 1) while the right image shows a kinematically possible alternative (hypothesis 2). If a linear motion model is used in the Kalman filter (constant velocity), we expect the tracker to choose hypothesis 2, which would be an incorrect description of the tracking scenario. With classification-aided tracking (CAT), however, the true motion shown in hypothesis 1 is correctly tracked 100% of the time.




FIGURE 26. Improvement in CAT performance as a function of prior knowledge. The target identification χ^2 values for correct and incorrect association are shown as a function of the *a priori* value. As prior knowledge increases, the gap between the χ^2 values increases, showing an improved discrimination between the correct and incorrect associations. As a result, the probability of correct classification also increases as the prior knowledge increases.

That is, as prior knowledge increases, the penalty for an incorrect association also increases, resulting in a better chance of correctly identifying the target. In addition, as the scan rate increases from ten seconds to five seconds, as expected, the probability of correctly identifying the target increases, with a very pronounced benefit for cases when the prior knowledge is large, such as seen when the prior equals 0.8.

Naturally, additional work is needed to quantify the impact of CAT upon tracker performance more reliably, but these preliminary results suggest that such additional work is warranted.

Two-Dimensional Feature-Aided Tracking Variants

Thus far we have considered the HRR profile to be the target feature, which is subsequently used for target identification or to help resolve kinematic ambiguities with the SAT and CAT feature-aided tracking algorithms. Although the HRR profile is readily obtainable with a high-bandwidth waveform, it does contain limited target feature information—a few tens of pixels represent the entire target. In addition, the HRR profile is known to be very sensitive to angle variation, thus significantly limiting the robustness of

HRR-based SAT algorithms. In contrast, a two-dimensional image of the target contains significantly more information about the target scattering. In addition, HRR profiles are available when the target is moving along a relatively straight path. When the target stops or turns, we have the opportunity to collect two-dimensional signatures. We show how these SAR and ISAR images may be used in a meaningful way to improve tracker performance.

In this section we show how two-dimensional images, in particular SAR and ISAR images, can be used in a two-dimensional variant of the SAT algorithm previously discussed. Furthermore, we show how the two-dimensional images are significantly more robust to aspect-angle variability. In addition, we recall that SAR and ISAR images can be used to generate HRR profiles at angles significantly different from those contained in the image; i.e., a few two-dimensional images can be used to populate an HRR template database covering a significant portion of the full 0°-to-360° range of aspect angles. This latter capability leads to the possibility of using SAR images of opportunity to construct an on-the-fly template set for a target type lacking previously collected templates for classification.

Two-Dimensional Signature-Aided Tracking

Consider a situation whereby a vehicle stops, starts moving again, and then stops. For a high-valued target, the operator could collect a SAR image of the target each time the target stopped. The degree of similarity between the two SAR images can be used to determine if the stopped vehicle is in fact the same vehicle in each image. The vehicles in the two images are, however, most likely at different aspect angles. Before a comparison of the SAR images can be made, the vehicles must be rotated to a common aspect angle, and before the vehicles can be rotated, an accurate estimate of the target aspect angle must be obtained. We have developed such a method but, for the sake of brevity, we do not include it in this discussion. Given the aligned, decluttered SAR images, the appropriately normalized two-dimensional signature χ^2 value is calculated from the MSE score obtained by comparing the two images. As with the one-dimensional version of SAT, where HRR profiles are compared, the signature χ^2 value is added to the kinematic χ^2 score to obtain a two-dimensional SAT score.

We anticipate that the MSE scores from aligned SAR images corresponding to the same vehicle (the match case) will be lower than the MSE score obtained from different vehicles (the mismatch case). Figure 27 shows the scores for the match and mismatch cases. The match scores were generated by comparing the SAR images of a Scud TEL at aspect angles separated by 30° , while the mismatch scores were obtained by comparing the SAR images of a Scud TEL and an M60 tank. It is evident that there is significant separation between the match and mismatch mean values, with minimal overlap in the distributions of scores. The fact that the two images are 30° apart is an indication that the SAR images are more robust to aspect-angle variability when used for SAT. However, part of the separation in distributions, possibly most of it, is due to the fact that the Scud TEL and M60 tank have different lengths. These results suggest that this two-dimensional SAT variant can provide additional benefit to the tracker. Future work will integrate this algorithm into the tracker, and we will perform further evaluation of the algorithm at that time.




FIGURE 27. The match and mismatch probability density functions resulting from comparing SAR images separated by 30° in aspect angle. For the match case, we compared SAR images of a Scud TEL separated by 30° in aspect angle and rotated so that the images were aligned in aspect angle. In the mismatch case, we compared SAR images of a Scud TEL against images of an M60 tank separated by 30° in aspect angle and rotated so the images were aligned in aspect angle. The MSE scores for the match case are significantly lower than the mismatch case, as expected, with minimal overlap in distribution.

We just described an algorithm that addressed the kinematic ambiguity that can result when a vehicle executes a “stop-move-stop” maneuver. In contrast, we now present a method for dealing with the so-called “move-turn-stop” scenario. Consider a vehicle that travels down a road, makes a turn, and then stops. If the operator has information about the road network, he can schedule an ISAR collection when the vehicle is anticipated to turn. If the operator is tracking a high-valued vehicle, he can justify expending the extra radar resources for a ISAR image since the two-dimensional image does contain more information about the target than an HRR profile. If the vehicle stops sometime thereafter, the operator can collect an SAR image of the area where the vehicle was last known to be located. The new variant of SAT that we are proposing compares the ISAR and SAR images in a manner similar to the procedure discussed for comparing two SAR images. As was described for comparing the SAR images, we must first align the ISAR and SAR images to a common aspect angle.




FIGURE 28. The match and mismatch probability density functions resulting from comparing SAR images of a Scud TEL against ISAR images of a Scud TEL and ISAR images of an M60 tank, respectively. The aspect-angle separation between the SAR and ISAR images was approximately 30° . For both match and mismatch cases, the SAR and ISAR images were aligned in aspect angle through rotation, and then compared. Moreover, the MSE computation required a search in cross range for the appropriate resolution.

Once aligned, a simple MSE calculation in range and cross range is performed, the appropriately normalized signature χ^2 value is calculated, and this value is added to the kinematic χ^2 value to obtain the total SAT χ^2 score. In Figure 28, we show the distribution of match and mismatch scores, where the match scores are obtained by comparing the SAR and ISAR images of a Scud TEL, which before alignment are separated by 30° , while the mismatch scores are obtained by comparing the SAR image of a Scud TEL with the ISAR image of a M60 tank, also initially separated by 30° . The separation in distributions between the match and mismatch scores provides further evidence that the target signatures present in the two-dimensional images are more robust to target aspect-angle variability than the HRR profile.

The last variant of two-dimensional SAT we present involves the comparison of two ISAR images. For this case, the operator has collected ISAR images at two different times when the vehicle is turning. With proper alignment of the targets to a common aspect angle, the MSE difference is calculated. In Fig-




FIGURE 29. The match and mismatch probability density functions as a result of comparing ISAR images of a Scud TEL against ISAR images of a Scud TEL and an M60 tank, respectively. The aspect-angle separation between ISAR images was approximately 30° . The MSE calculation between a pair of ISAR images required a search for the appropriate cross-range resolution for both images.

ure 29, the distribution of match and mismatch scores is calculated, as was done for the previous cases in Figure 27 (SAR versus SAR) and Figure 28 (SAR versus ISAR). We see reasonable separation between the match and mismatch scores, thus providing the basis for hope that, with two-dimensional SAT, the tracker can make beneficial use of the extra information present in the images.

Two-Dimensional Classification-Aided Tracking

CAT uses target identification to help the tracker resolve kinematically ambiguous situations. We have presented a method for obtaining an estimate of the target identification by using HRR profiles. In this section, we show how target identification obtained from two-dimensional images—SAR and ISAR—can also be used. This estimate of target identification can be used to obtain a χ_{ID}^2 score that can be added to the kinematic score to obtain the CAT χ^2 value.

When a high-valued target stops, it is possible to collect a SAR image of the area where the vehicle was last known to be moving. A chip of the target from the SAR image can be obtained by using fixed-target

indication and extraction methods. Given the target chip, the target identification can be estimated by using a template-based matching scheme. On the basis of the Semi-Automated Image Intelligence Processing (SAIP) work of Novak [2], we can expect the estimate of the target identification obtained with the SAR to be significantly better than that obtained with an HRR profile. For example, in Figure 30, we show the SAR image and HRR profile of the Scud TEL. Clearly, more information about the target structure is present in the two-dimensional SAR images than in the one-dimensional HRR profiles. Furthermore, with the advances we have made in superresolution techniques—BHDI in particular—the estimate of the target identification is better than previously reported. In Figure 31, we show the gain in performance, using a SAIP ATR classifier, for a six-class classifier, as measured by a ROC curve in BHDI ver-

sus baseline HDVI and conventional baseline FFT processing. The improvement in performance, especially at the operationally significant lower Pfc values, say 0.1 to 0.01, is significant.

When a high-valued target turns, an operator may collect an ISAR image. A template-based matching scheme may also be used for ISAR images. In this case, however, we need to compare ISAR images against SAR images in the template database. Figure 32 illustrates an example of this type of comparison. Previous work by Levin on template-based ISAR ATR [4], and our work on signature-aided tracking, summarized in Figure 28, showed that there are sufficient differences between Scud TEL and M60 tank ISAR images that reasonable discrimination can be obtained. Our ongoing feature-aided tracking work incorporates the ISAR mode into the tracker to help improve both SAT and CAT.




FIGURE 30. The two upper images depict a SAR image of a Scud TEL and its best-matched SAR template, while the graphs below show HRR profiles corresponding to the test and template images. Visual inspection of these images and graphs clearly signifies that the two-dimensional SAR mode provides significantly higher matching confidence than the one-dimensional HRR counterpart.




FIGURE 31. Performance results using the Semi-Automated Image Intelligence Processing (SAIP) two-dimensional ATR classifier for SAR images of the same six vehicles demonstrated in the HRR case (and shown in Figure 7). These results indicate that BHDl-processed SAR images provide the best ATR performance, compared to baseline HDVI and Taylor-weighted FFT methods. More importantly, the two-dimensional classifier using SAR images provides much more reliable classification confidence than the one-dimensional classifier using HRR profiles.

Summary

Identification and tracking of moving targets are coupled problems, consisting of identification tech-

niques that use successive target signatures and accumulate classification information in a tracker, and feature-aided tracking techniques that take advantage of target signatures and accumulated classification information to improve tracker performance. In this article we consider the problems of improving moving-target classification and track continuity by using radar signatures of the targets, e.g., HRR, SAR, and ISAR. Our feature-aided tracker incorporates these radar signatures into a kinematic tracker algorithm [16] using two modes: classification and verification.

In the classification mode, a one-dimensional template-based automatic target recognizer provides target identification based on the target's HRR profile. For operational consideration, we used feature-enhancement algorithms such as Benitz's BHDl to improve the feature quality of the HRR profiles. We demonstrate a Bayesian classifier using the resulting super-resolved profiles to provide moving-target identification. By incorporating the classifier into the traditional tracker, we then established a technique by which the HRR feature information can be fused with the kinematic information to improve track-to-report association and resolve kinematic ambiguities. To demonstrate the performance of the classification-aided mode, we simulated a kinematically confusing two-target scenario by emulating target MTI detec-




FIGURE 32. Comparison of an ISAR image and a SAR template image of the same target, showing the degree of similarity of the images. The ISAR image on the left shows a Scud TEL moving in a circular arc of 50-m radius and turning at a rate of 9.78° per second. The ISAR image has a 1-ft range resolution and approximately 1-ft cross-range resolution. The SAR image on the right shows the exact same target, also at 1-ft by 1-ft resolution.

tions augmented with HRR profiles by merging noise-injected GPS data with profiles generated from SAR images taken from the MSTAR and Data Collection system (DCS) data sets. The CAT mode successfully resolved all tracks-reports misassociations made by the kinematic tracker. Furthermore, we showed that the performance improvement provided by ATR can be quantified by *a priori* knowledge of the tracks' vehicle types.

In the verification mode, track and report signatures are systematically compared, presuming that these signatures are similar if they come from the same target. We demonstrated an approach to fuse the signature comparison results with the UAV's kinematic information by using a log-likelihood ratio. Applying this approach to a simulated kinematically confusing two-target scenario, we showed that the signature-aided component provided modest improvement to the traditional tracker due to HRR feature variability on aspect angle.

To improve the performance of the SAT mode, we considered a supplemental approach to handle large aspect differences by incorporating SAR and ISAR images into the tracker. We demonstrated metrics to compare HRR, SAR, and ISAR signatures, and we extended the range of aspect-angle mismatch seen in the SAT mode with only HRR profiles. When we encountered track and report profiles whose aspect angles differed by more than 10° , we used a two-dimensional image associated with the track to generate an HRR profile at the same aspect angle as that of the report profile. To obtain SAR images while tracking, we considered a go-stop-go scenario in which the targets stopped long enough for the radar to obtain a SAR image. To obtain ISAR images while tracking, we assumed the availability of a road network to predict ISAR opportunities.

The classification and verification modes of the feature-aided tracker can potentially operate simultaneously to provide improved track continuity and target identification.

Acknowledgments

The authors would like to thank Gerald Benitz of the Sensor Exploitation group at Lincoln Laboratory for providing us with his expertise on superresolution

techniques, as well as his HDVI and BHDI codes. We would also like to thank Keith Sisterson, also of the Sensor Exploitation group, for providing us with the UAV tracker. This work was sponsored by DARPA.

REFERENCES

1. M.I. Mirkin, B.E. Hodges, and J.C. Henry, "Moving Target Recognition," presentation for DARPA ATR Exposition, 17 Sept. 1998.
2. L.M. Novak, G.J. Owirka, W.S. Brower, and A.L. Weaver, "The Automatic Target-Recognition System in SAIP," *Linc. Lab. J.* **10** (2), 1997, pp. 187–202.
3. G.R. Benitz, "High-Definition Vector Imaging," *Linc. Lab. J.* **10** (2), 1997, pp. 147–170.
4. R.L. Levin, private conversation.
5. D.H. Nguyen, G.R. Benitz, J.H. Kay, and R.H. Whiting, "Super-Resolution HRR ATR Performance with HDVI," *SPIE* **4050**, 2000, pp. 418–427.
6. D.H. Nguyen, G.R. Benitz, J.H. Kay, B.J. Orchard, and R.H. Whiting, "Super-Resolution High Range Resolution ATR with HDVI," *IEEE Trans. Aerosp. Electron. Syst.* **37** (4), 2001, pp. 1267–1286.
7. D.H. Nguyen, J.H. Kay, B.J. Orchard, and R.H. Whiting, "Improving HRR ATR Performance at Low SNR by Multi-Look Adaptive Weighting," *SPIE* **4379**, 2001, pp. 216–228.
8. D.H. Nguyen, G.R. Benitz, J.H. Kay, B.J. Orchard, and R.H. Whiting, "HRR ATR Performance Enhancement at Low SNR Using Beamspace HDI," *SPIE* **4379**, 2001, pp. 266–276.
9. G.J. Owirka, S.M. Verbout, and L.M. Novak, "Template-Based SAR ATR Performance Using Different Image Enhancement Techniques," *SPIE* **3721**, 1999, pp. 302–319.
10. D.H. Johnson, "The Application of Spectral Estimation Methods to Bearing Estimation Problems," *Proc. IEEE* **70** (9), 1982, pp. 1018–1028.
11. H.C. Stankwitz, R.J. Dallaire, and J.R. Fienup, "Non-Linear Apodization for Sidelobe Control in SAR Imagery," *IEEE Trans. Aerosp. Electron. Syst.* **3** (1), 1995, pp. 267–279.
12. H.C. Stankwitz and M.R. Kosek, "Super-Resolution for SAR/ISAR RCS Measurement Using Spatially Variant Apodization (Super SVA)," *Proc. AMTA Symp., Williamsburg, Va., 13–7 Nov. 1995*, pp. 251–256.
13. J. Capon, "High-Resolution Frequency-Wavenumber Spectrum Analysis," *Proc. IEEE* **57** (8), 1969, pp. 1408–1418.
14. G.R. Benitz, "Beamspace High-Definition Imaging (B-HDI) SAR Image Enhancement," submitted to *IEEE Trans. Image Process.*
15. J.H. Kay and R. Levin, internal presentation.
16. L.K. Sisterson, private conversation.
17. D.H. Nguyen, J.H. Kay, B.J. Orchard, and R.H. Whiting, "Classification-Aided Tracking," *Proc. 55th ATRWG Science and Technology Symp., Aerospace Corp., Chantilly, Va., 6–18 July 2001*.
18. D.H. Nguyen, J.H. Kay, B.J. Orchard, and R.H. Whiting, "Feature-Aided Tracking of Ground Moving Vehicles," *SPIE* **18**, 2002, pp. 234–245.

APPENDIX: BEAMSPACE HIGH-DEFINITION IMAGING SUPERRESOLUTION METHOD

WE PRESENT A BRIEF DISCUSSION of the new superresolution method called beamspace high-definition imaging (BHDI), as applied to synthetic aperture radar (SAR) images and high range resolution (HRR) profiles. This method was originally developed by Gerald R. Benitz of Lincoln Laboratory for forming SAR images. Since BHDI is related to high-definition vector imaging (HDVI), we first discuss HDVI, and how it is related to J. Capon's method, and then we show how HDVI is extended to BHDI.

Recently Benitz proposed a modification to HDVI that reduces the computational cost of BHDI to one-sixth the amount needed for HDVI. In particular, HDVI required 8000 operations per pixel. With BHDI, only 8000/6 operations are needed. BHDI provides speckle noise suppression that is comparable to HDVI but improves and better preserves target edges. In Figure 1 in the article we compare a BHDI-processed SAR image of a Scud TEL with an image formed with the conventional Taylor-weighted FFT (baseline FFT) method. The inherent ability of BHDI to control the main lobe and the sidelobes is clearly visible as evidenced by the significant reduction in speckle noise in the BHDI image compared to the weighted FFT.

With this success with BHDI in forming SAR images, we adapted the method to the one-dimensional case for forming HRR profiles. The superresolution of HRR data can be viewed as a parameter estimation problem in which we seek to estimate the reflectivity intensity at each pixel by selecting an appropriate set of weighting coefficients. The basic one-dimensional spectral estimation problem can be posed in the form of Capon's maximum likelihood method (MLM), by which it is desired to estimate the power spectrum $\rho(\cdot)$ such that

$$\begin{aligned}\rho(x) &= \min_{\mathbf{w}} E \left\{ \left\| \mathbf{w}^H \mathbf{r} \right\|^2 \right\} \\ &= \min_{\mathbf{w}} [\mathbf{R} \mathbf{w}, \mathbf{w}],\end{aligned}$$

subject to the constraint

$$[\mathbf{v}, \mathbf{w}] = 1,$$

where $\mathbf{w} \equiv \mathbf{w}(x)$ is a weight vector, $\mathbf{r} \equiv \mathbf{r}(x)$ is an HRR profile, $\mathbf{v} \equiv \mathbf{v}(x)$ is an ideal point scatterer response, and $\mathbf{R} \equiv \mathbf{R}(x)$ is a covariance matrix at pixel location x . We used $E\{\cdot\}$ to denote the expectation, $\|\cdot\|$ to denote the Euclidean L_2 -norm, and $[\cdot, \cdot]$ to denote the vector inner-product operator. For an invertible covariance matrix $\mathbf{R}$, the spectral estimate $\rho(x)$ is given by the well-known result

$$\rho(x) = \frac{1}{[\mathbf{R}^{-1} \mathbf{v}, \mathbf{v}]}.$$

Since $\mathbf{R}$ is unknown, it must be estimated before Capon's technique can be applied. Several strategies exist for estimating a covariance matrix, including subband averaging, forward-backward subband averaging, and block-Toeplitz methods. In our work, we employed the forward-backward subband averaging technique in which each HRR is sectioned into overlapping subsets whose correlations are subsequently averaged to obtain an estimated covariance matrix $\hat{\mathbf{R}}$. Let $H = \{\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_M\}$, where $\mathbf{r}_i \in R^N (i = 1, \dots, M)$, denotes a set whose elements represent N -sample HRR profiles. If q subbands from each profile are used to estimate the covariance matrix, then $\hat{\mathbf{R}}$ is given by

$$\hat{\mathbf{R}} = \frac{1}{2} (\tilde{\mathbf{R}} + \mathbf{J} \tilde{\mathbf{R}} \mathbf{J}),$$




FIGURE A. Comparison of HRR profiles of a T72 tank formed by using the weighted FFT, HDVI, and BHDl algorithms. The graph at the bottom left shows an HRR profile at 0.5-m range resolution and 20-dB SNR, while the other three graphs show the same data at 1-m range resolution. The HDVI-processed profile at 1-m range resolution clearly exhibits scattering responses resembling those of the 0.5-m range-resolution profile processed with the weighted FFT. However, HDVI failed to preserve both of the scattering responses at the edges of the T72. BHDl, on the other hand, successfully preserved all the scattering responses seen in the HDVI-processed profile as well as both of the responses at the T72's edges.

where

$$\tilde{\mathbf{R}} = \frac{1}{Mq} \sum_{i=1}^{Mq} [\mathbf{r}_i, \mathbf{r}_i]$$

and

$$\mathbf{J} = \begin{bmatrix} 0 & 0 & \cdots & 0 & 1 \\ 0 & 0 & \cdots & 1 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 1 & \cdots & \vdots & 0 \\ 1 & 0 & \cdots & 0 & 0 \end{bmatrix}_{(N-q+1) \times (N-q+1)}.$$

When the estimated covariance matrix is singular, HDVI introduced a twofold modification to Capon's

technique to accommodate a noninvertible $\hat{\mathbf{R}}$. The first modification, called the *quadratic constraint*, limits the norm of the weight vector from exceeding some threshold. The second modification, called the *subspace constraint*, constrains the weight vector to the subspace spanned by $\hat{\mathbf{R}}$.

The basic idea of HDVI is to solve the constrained optimization problem subject further to

$$\|\mathbf{w} - \mathbf{w}_{\text{proj}}\|^2 \leq \beta$$

and

$$[\mathbf{v}_{\text{proj}}, \mathbf{w}] = \|\mathbf{v}_{\text{proj}}\|^2,$$

where $\mathbf{w} \in \text{span}\{\hat{\mathbf{R}}\}$. Here, $\mathbf{v}_{\text{proj}}$ denotes the projec-

tion of $\mathbf{v}$ onto the column space of $\hat{\mathbf{R}}$, and β is a positive real number used to limit how far away $\mathbf{w}$ can deviate from $\mathbf{v}_{\text{proj}}$.

BHDI employs a different approach to estimate the radar cross section at each pixel of interest. The basic assumption in BHDI is that most of the energy contamination to the pixel of interest comes mainly from the adjacent pixels. To estimate the radar cross section at each pixel of interest, we need only to add and subtract the contributions from the adjacent pixels. Let $\mathbf{r}_{ij}$ denote the j th pixel of the i th profile $\mathbf{r}_i$, and define a vector $\mathbf{z}_{ij}$ as follows:

$$\mathbf{z}_{ij} = \begin{bmatrix} \mathbf{r}_{ij} \\ \mathbf{r}_{i(j-1)} + \mathbf{r}_{i(j+1)} \end{bmatrix}.$$

The estimated covariance matrices of the real and imaginary components at the j th pixel of the i th profile hence depend only on the adjacent pixels ($j - 1$) and ($j + 1$):

$$\left. \begin{aligned} \hat{\mathbf{R}}_{ij}^{\text{re}} &= \mathbf{z}_{ij}^{\text{re}} (\mathbf{z}_{ij}^{\text{re}})^H \\ \hat{\mathbf{R}}_{ij}^{\text{im}} &= \mathbf{z}_{ij}^{\text{im}} (\mathbf{z}_{ij}^{\text{im}})^H \end{aligned} \right\} \quad (2 \times 2 \text{ matrices})$$

The estimated power spectrum $\rho(j)$ at pixel j is thus given by the equality

$$\rho(j) = \min_{\mathbf{w}_j} \left\{ \mathbf{w}_j^H \underbrace{\left(\frac{1}{M} \sum_{i=1}^M (\hat{\mathbf{R}}_{ij}^{\text{re}} + \hat{\mathbf{R}}_{ij}^{\text{im}}) \right)}_{\hat{\mathbf{R}}_j} \mathbf{w}_j \right\},$$

where

$$\mathbf{w}_j = \begin{bmatrix} 1 \\ \xi \end{bmatrix}$$

and $\xi \in \mathbb{R}$, for $j = 1, \dots, N$, where N is the number of range pixels in each profile. We note that as long as $M > 1$, the estimated covariance matrix at pixel j , denoted $\hat{\mathbf{R}}_j$, will be nonsingular. As an effect, the subspace constraint developed originally for HDVI is no longer required.

We now discuss a sample result using BHDI. Figure A shows sample HRR profiles of a T72 tank at 20-dB SNR and 1-m and 0.5-m range resolutions. The baseline-FFT-processed profile at 1-m resolution (upper left in the figure) exhibits two broad scatterer responses, whereas the baseline HDVI-processed profile (upper right) reveals two additional scatterer responses. When we compare the 0.5-m range-resolution profile processed with the baseline FFT (lower left), we see that the additional scatterer responses shown in the baseline HDVI-processed profile also appear at the higher resolution. However, the baseline HDVI profile failed to detect the scatterer response that appears along the left edge of the 0.5-m resolution baseline FFT profile. The application of BHDI (lower right) successfully picked up the scatterer responses at both edges of the target. Since the baseline HDVI algorithm attempts to estimate the radar cross section at each pixel of interest by adding and subtracting out contributions from every other pixel, the algorithm may be treating noisy information along the edges as relevant information. BHDI, however, uses only the adjacent pixels and is thus less prone to contributions from noisy scatterer responses about the edges. As an effect, BHDI tends to preserve the scatterer responses toward the edges of the target, while HDVI is inclined to null them out.



DUY H. NGUYEN

was born in Saigon, Vietnam, in 1971. He completed his course work for the mathematics degree from the University of California, San Diego, in 1989, received a B.S. degree in mathematics and engineering from Harvey Mudd College, Claremont, California, in 1993, and M.S. and Ph.D. degrees in electrical engineering from the University of California, Los Angeles, in 1995 and 1998, respectively. From 1994 to 1997 he was a member of the technical staff in the control analysis group at the Jet Propulsion Laboratory. From 1997 to 1998 he worked as a systems engineer in the TRW Avionics Systems Division in San Diego, California. In 1999, Dr. Nguyen joined Lincoln Laboratory as a member of the technical staff in the Signals and Systems group. His current interests include feature-aided tracking, automated target recognition, superresolution, complex systems, and self-organizing control.



JOHN H. KAY

received a B.S. degree from Purdue university in 1974 and a Ph.D. degree from Brandeis University in 1978, both in physics. He has worked at Lincoln Laboratory since 1978, where he is a senior staff member, currently working in the Signals and Systems group. His research interests include automatic target recognition as it applies to both moving and stationary targets.



BRADLEY J. ORCHARD

is a staff member in the Signals and Systems group. He received a B.S. degree from Loyola University at Chicago, an M.S. degree from the University of Illinois at Chicago, and a Ph.D. degree from Rensselaer Polytechnic Institute, all in mathematics. In addition, he has undergraduate and graduate degrees in chemistry and polymer science. Prior to joining Lincoln Laboratory, he worked at TASC in Reading, Massachusetts, and at the U.S. Naval Research Laboratory in Washington, D.C. He was also an adjunct assistant professor of mathematics and computer science at Merrimack College and Bentley College. He has been involved in research in a variety of fields, including solid state physics, underwater acoustics, parallel processing, signal processing, multisensor data fusion, and moving-target detection, tracking, classification, and imaging.



ROBERT H. WHITING

received B.S., M.S., and Ph.D. degrees in electrical engineering at the University of Idaho, Stanford University, and the University of Idaho, respectively. His thesis work was on the study of optimization techniques for the control of sequential hydroelectric plants to maximize power output with environmental planning constraints. He joined Lincoln Laboratory in 1973, and he is currently group leader with the Signals and System group. His areas of research include radar technology, application of estimation and optimization theory, remote sensing and target recognition techniques, and image processing. He is a member of the IEEE.