SAND2007-1307C

Multinomial Pattern Matching for High Range Resolution Radar Profiles

Melissa L. Koudelka, John A. Richards, and Mark W. Koch

Sensor Exploitation Applications
Sandia National Laboratories
Albuquerque, New Mexico, USA

ABSTRACT

Airborne ground moving-target indication (GMTI) radar can track moving vehicles at large standoff distances. Unfortunately, trajectories from multiple vehicles can become kinematically ambiguous, resulting in confusion between a target vehicle of interest and other vehicles. We propose the use of high range resolution (HRR) radar profiles and multinomial pattern matching (MPM) for target fingerprinting and track stitching to overcome kinematic ambiguities.

Sandia's MPM algorithm is a robust template-based identification algorithm that has been applied successfully to various target recognition problems. MPM utilizes a quantile transformation to map target intensity samples to a small number of grayscale values, or quantiles. The algorithm relies on a statistical characterization of the multinomial distribution of the sample-by-sample intensity values for target profiles. The quantile transformation and statistical characterization procedures are extremely well suited to a robust representation of targets for HRR profiles: they are invariant to sensor calibration, robust to target signature variations, and lend themselves to efficient matching algorithms.

In typical HRR tracking applications, target fingerprints must be initiated on the fly from a limited number of HRR profiles. Data may accumulate indefinitely as vehicles are tracked, and their templates must be continually updated without becoming unbounded in size or complexity. To address this need, an incrementally updated version of MPM has been developed. This implementation of MPM incorporates individual HRR profiles as they become available, and fuses data from multiple aspect angles for a given target to aid in track stitching. This paper provides a description of the incrementally updated version of MPM.

Keywords: Multinomial pattern matching, HRR, target fingerprinting, target tracking, track stitching

1. INTRODUCTION

Moving-vehicle tracking is an application of considerable military importance and interest. Radar provides an attractive modality for the implementation of effective target tracking applications: it offers an all-weather, day/night, long-standoff capability for the detection and tracking of moving vehicles. Ground moving-target indication (GMTI) radar systems have been developed to enable effective target detection and tracking.[1] GMTI-based systems are typically quite successful in detecting and tracking particular targets of interest when they are operating in the absence of other moving vehicles. In the presence of other vehicles, however, GMTI-based tracking applications are susceptible to loss of track due to kinematic ambiguities between the trajectories of targets of interest and other vehicles. Because only extremely limited vehicle-specific information is provided by GMTI radar, GMTI imaging modes are often implemented in conjunction with high-range-resolution (HRR) imaging modes.[2] HRR imaging enables the collection of HRR profiles detailing range-dependent target returns, or *profiles*, which differ significantly from vehicle to vehicle. HRR profiles provide vehicle-specific signatures that can be used to resolve kinematic ambiguities and extend track life. In particular, HRR profiles can be used to develop target fingerprinting applications that facilitate the recognition of individual vehicles, thus enabling GMTI track maintenance in the presence of kinematic ambiguities.

A target fingerprinter is typically developed and implemented for use in conjunction with a number of other radar subsystems, including a sensor resource manager (SRM) and kinematic tracker (KT).[3] In general, each subsystem is responsible for a specific set of tasks that, together, would provide a complete solution to the problem of acquiring and maintaining tracks on vehicles of interest. The KT is responsible for grouping individual GMTI returns into kinematically unambiguous *tracklets* representing trajectories of individual vehicles of interest, for determining when tracklets become, or are in danger of becoming, kinematically ambiguous with other tracklets, and for stitching together tracklets that have been determined to correspond to the same target. The SRM is

Corresponding author information: M. L. Koudelka, mlkoude@sandia.gov, (505) 284-8843.

responsible for allocation of radar resources, including radar pointing and image mode selection. Its resource allocation is based at least partially on information and requests from the KT: if the KT determined that two vehicle tracklets were in danger of becoming kinematically ambiguous, for instance, it might submit a request to the SRM to collect HRR data on both vehicles in order to facilitate their disambiguation by the target fingerprinter, and subsequent restitching by the KT. The target fingerprinter is responsible for encapsulating HRR signatures from each distinct target tracklet into target fingerprints, and for calculating match scores and confidences corresponding to specific target-association hypotheses when requested to do so by the KT in order to resolve kinematic ambiguities.

In practice, HRR-based target fingerprinting is a difficult problem. HRR profiles are essentially coherent integrations of target reflectivity as a function of range. As such, they tend to vary significantly with small changes in target aspect, and to exhibit significant variations in peak amplitudes due to constructive and destructive interference effects. These effects make the development and implementation of a statistical framework for the characterization and exploitation of HRR target signatures a nontrivial matter. Sandia's multinomial pattern matching (MPM) algorithm addresses the difficult issues associated with HRR-based target fingerprinting. Section 2 provides a description of the motivating principles of MPM. Section 3 details the process by which individual HRR profiles are used to incrementally develop MPM fingerprints. The MPM scoring process is described in Section 4. Experimental results illustrating the utility of MPM for target fingerprinting are presented in Section 5. Section 6 summarizes the paper.

2. MPM MOTIVATION

Sandia's MPM algorithm was originally developed for use as an automatic target recognition (ATR) identification algorithm[4] for use on synthetic aperture radar (SAR) imagery.[5] It is designed to address many of the fundamental characteristics of radar signatures—whether SAR or HRR—that make robust recognition of these signatures so challenging. In particular, it is robust to target amplitude variation of the type present in SAR, HRR, and other coherent imaging modalities; it is effective in the presence of limited target signature contamination and variation; it is robust to sensor calibration.

MPM provides a robust means for statistical characterization of radar signatures and for manipulation of those characterizations to yield match scores that can serve as the basis for target identity declarations. MPM is motivated by three fundamental premises:

1. *Intra-class modeling premise.* Target identity declarations (and, by extension, the models used to obtain declarations) should be based on intra-class similarities—that is, on specific characteristics inherent to a particular target class—and *not* on differences between target classes. Target models that characterize individual classes can be used directly to reject signatures that do not conform to a previously observed target type, because such models are designed to classify signatures as "target A" or "not target A". Target models based on the *differences* between classes, however, cannot be used directly to reject signatures from previously unobserved targets: they are designed to classify signatures as "target A" or "target B".

2. *Signature stability premise.* Target models should be based on the stable information inherent in the signatures of a given class, and should, as much as possible, discard highly variant or nuisance information within each class. In the context of HRR-based fingerprinting, this suggests that target models should be based on information such as the relative locations of bright and dim samples (which are stable within a target class over moderate changes in aspect) instead of on the precise amplitudes of samples in these profiles (which are extremely fickle and highly variant with small changes in aspect).

3. *Match robustness premise.* Target identity declarations should be based on match statistics that are robust to isolated outlying measurements. For instance, individual arbitrarily bad sample values within a profile should not be able to drive association results, but should have a limited impact on the overall match statistic and resulting association decision.

These three premises suggest a particular form for an effective target fingerprinting algorithm, and serve as the basis for the MPM algorithm. A high-level block diagram of MPM is depicted in Figure 1. Target profiles from kinematically unambiguous tracklets are separated into distinct sets, and then subjected to a fingerprinting process that is independent of the profiles associated with other tracklets by the KT, as motivated by the intra-class modeling premise. The details of the statistical characterization inherent in the fingerprinting process are motivated by the signature stability premise, and are structured to yield a model that captures the relevant, stable signature information within each target class. Fingerprints of distinct sets of tracklets—for instance, one set of pre-ambiguity tracklets preceding a period of kinematic confusion, and a second set of post-ambiguity

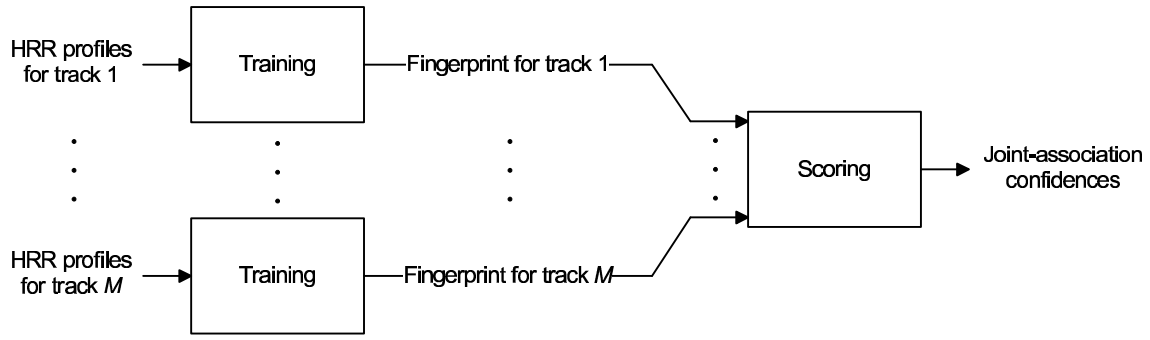

Figure 1. Top-level block diagram of MPM training and tracklet-association scoring.

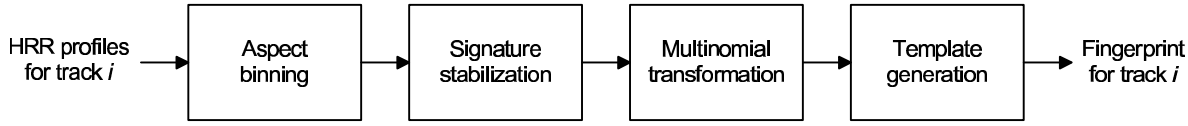

Figure 2. Block diagram of MPM training process.

tracklets following a period of kinematic confusion—can be compared using a match statistic motivated by the match robustness premise to yield a set of match scores which can in turn be used to generate a confidence matrix specifying levels of belief in each possible resolution of a particular kinematic ambiguity.

The algorithmic details and implementation of the training component of MPM is described in Section 3. The algorithmic details of the scoring component of MPM is described in Section 4.

3. MPM TRAINING DETAILS AND IMPLEMENTATION

MPM training is a continuous incremental process, in which additional profiles are incorporated into existing tracklet fingerprints as soon as they become available. For any given tracklet, MPM maintains a set of multiple fingerprints, each representing a collection of HRR profiles in a particular aspect range. An MPM fingerprint, or *template*, consists of two components: first, a statistical characterization of the HRR profiles for the particular tracklet and aspect range, and second, a set of sample penalties that enable the fingerprint to be scored against any other HRR fingerprint or observed HRR signature.

A block diagram of the MPM training process is depicted in Figure 2. (Note that each individual tracklet is in effect an independent training stream, so that if the overall system is maintaining track on M separate vehicles, there will be M separate parallel training streams.) Profiles within a tracklet are first binned by aspect using aspect estimates provided by the KT. (Typically, the MPM aspect bin width is chosen to be $10°$, which is wide enough to populate aspect bins with multiple profiles over even a short tracking engagement, but narrow enough to preclude drastic changes in target signature over the width of the aspect bin.) The profiles within each aspect bin are then stabilized to increase self-similarity, as described in Section 3.1. This is followed by two key steps in the MPM training process: first, multinomial transformation, described in Section 3.2, and second, generation of the actual MPM template, described in Section 3.3.

3.1. Profile Stabilization

Profile stabilization consists of several steps intended to maximize the self-similarity between profiles within any aspect bin. Maximizing this self-similarity is an important prerequisite for the ensuing training steps, because in effect, it removes nuisance information and noise that would significantly complicate or invalidate the following stages.

MPM profile stabilization consists of profile alignment followed by optional length-normalization and smoothing steps. Profile alignment is performed using a leading-edge detector that estimates the location of the nearest-range target sample in a profile. All profiles in a given aspect bin are cropped or zero-padded so that their estimated leading edges fall into the same range bin, and so that their overall lengths are identical.

For narrow aspect bins, this alignment process generally yields an acceptable registration between profiles. For wider aspect bins, in which the apparent length of the target may change appreciably across the aspect bin,

target length normalization may aid registration. If desired, this length normalization is implemented using a trailing-edge detector to estimate target length in each profile, and then interpolating to stretch or compress the target component of the profile to a fixed number of samples.

Profile smoothing is an optional final step in the stabilization process. Smoothing involves a fundamental trade-off between resolution and stability. Unsmoothed data will display more fine-resolution structure, including fine-resolution features that might aid target characterization; at the same time, unsmoothed data will exhibit more noise and spurious amplitude variations that are not truly informative. The use or omission of the optional smoothing step is dictated by the particular characteristics of the data under consideration.

3.2. Multinomial Transformation

A key step in the MPM training process is multinomial, or quantile, transformation. This transform maps each sample in a profile from a raw magnitude to a quantile index specifying the relative intensity of that sample in comparison to all other profile samples. It thus transforms profiles from representations of absolute amplitude to representations of relative amplitude. This quantile transform is motivated by the second premise of Section 2: it preserves information that is stable within a target class, and is easy to model—namely, the locations of bright and dim samples—and discards information that is highly variant within a target class, and is relatively difficult to model—namely, the nature of the precise amplitude variations of individual samples. Rather than attempting to model a continuous amplitude distribution at each sample, MPM attempts to model a discrete quantile distribution—typically a much easier task.

The MPM quantile transform is defined as follows. Let x_k be the observed magnitude of sample k of an HRR profile, where k is a sample index that takes on integer values between 1 and K. The K profile samples can be rank-ordered by magnitude and assigned rank-order indices $r_1, \ldots, r_K$ so that

$$x_{r_1} \leq x_{r_2} \leq \ldots \leq x_{r_K}. \tag{1}$$

The quantile transform that maps magnitude values x_k to quantile values q_k is then defined in terms of the rank-order indices and a specified number of quantiles, N_q, as follows:

$$q_k = \begin{cases} 1 & \text{if} \quad 1 \leq r_k \leq \frac{K}{N_q}, \\ 2 & \text{if} \quad \frac{K}{N_q} < r_k \leq \frac{2K}{N_q}, \\ \vdots & \\ N_q & \text{if} \quad \frac{(N_q-1)K}{N_q} < r_k \leq K. \end{cases} \tag{2}$$

The dimmest K/N_q profile samples in a profile thus are mapped to quantile 1, the next-dimmest K/N_q samples are mapped to quantile 2, and so on, with the brightest K/N_q samples mapped to quantile N_q. An illustration of the quantile transform is depicted in Figure 3.

Although there is certainly more information in the collection of x_k than in q_k, the information discarded by the quantile mapping—namely, the precise nature of amplitude variations at any sample—is extremely variable and difficult to model reliably with a limited amount of data. The information preserved by the quantile mapping—namely, the *relative* amplitude at each sample, compared to other samples in the same profile—is much more stable and much better suited for efficient, robust statistical modeling.

3.3. MPM Template Generation

The quantile-transformed profiles within each aspect bin are used to generate an MPM template for each aspect bin. As described previously, an MPM template consists of two components: first, a compact statistical characterization of the quantized profiles from the tracklet and aspect bin in question, and second, a compact collection of sample penalties that enable templates to be scored against each other, as described further in Section 4. The construction of the statistical-characterization component of an MPM template is described in Section 3.3.1; the construction of the sample-penalty component is described in Section 3.3.2.

3.3.1. MPM Template: Statistical-Characterization Component

In principle, a set of HRR profiles could be "characterized" simply by storing all members of the set. In practice, however, such an exhaustive characterization would be cumbersome or impossible. Over the course of a single target engagement, many tens of thousands of HRR profiles might be collected. Storing a complete collection of profiles, and manipulating these profiles to yield match scores in real time, would be impractical or impossible. In practice, it is necessary to encapsulate the relevant statistical features of profile sets into a compact, fixed-size representation that lends itself to efficient manipulation.

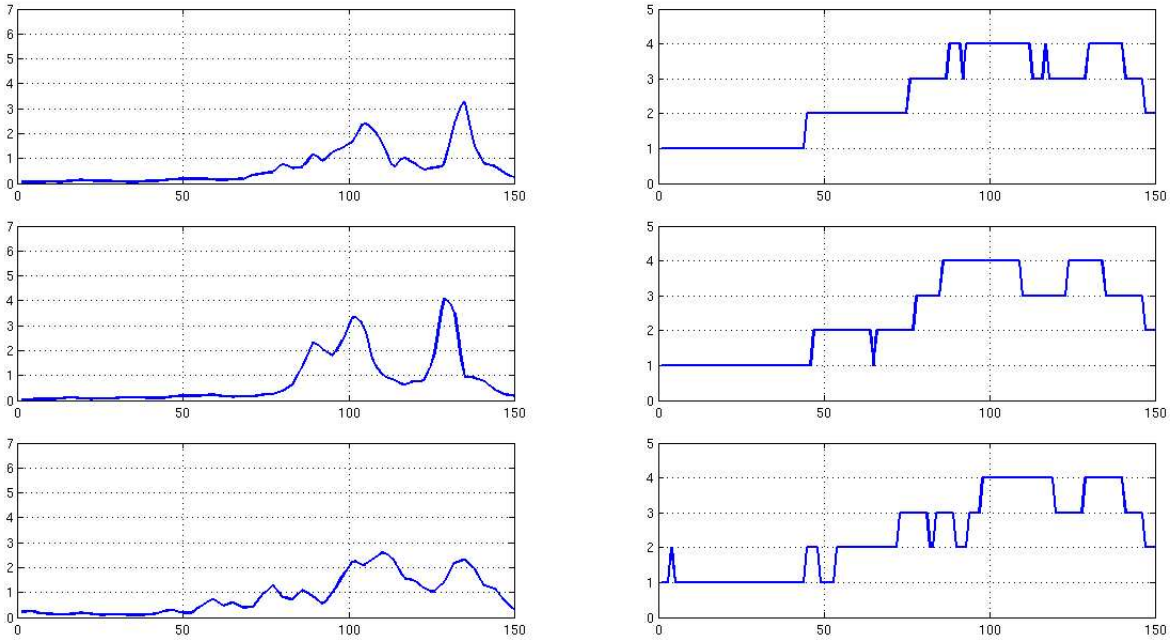

Figure 3. Illustration of multinomial transformation. The three plots on the left are raw-magnitude HRR profiles; the three plots on the right are the corresponding HRR profiles after multinomial transformation.

Quantile transformation produces data that is perfectly suited to compact representation and to robust statistical estimation given a limited amount of data. In particular, the probability distribution of an N_q-quantile-level sample is completely specified by $N_q - 1$ parameters (the N_qth quantile probability is constrained to make the collection of N_q quantile probabilities sum to 1), and can be robustly estimated with relatively few (*i.e.*, not many more than $N_q - 1$) observations. Hence, if all profiles are cropped or padded to a fixed number of samples K, then the collection of marginal sample distributions of a set of quantile-transformed HRR profiles can be specified with only $(N_q - 1)K$ parameters, and robustly estimated with not many more than $N_q - 1$ profile observations. (Of course, a marginal sample-distribution model neglects dependencies between observed quantile values at different sample locations, but robust estimation of such parameters would require significantly more data than robust estimation of the simpler marginal model.)

In practice, it is intuitive and convenient to specify the marginal quantile model as a $K \times N_q$ matrix of observed quantile probabilities $\hat{\mathbf{P}}$, with the row-k, column-m entry $\hat{P}_{k,m}$ corresponding to the empirical probability of the kth sample having quantile value m. This empirical-probability matrix not only provides a complete marginal specification of the empirical sample quantile probabilities, but it has extremely convenient computational properties: it is compact and fixed size, and it lends itself perfectly to incremental updates as additional profiles become available within a class. In particular, if the empirical-probability matrix after N profile observations is specified by a matrix $\hat{\mathbf{P}}[N]$ of values $\hat{P}_{k,m}[N]$, then elements of the updated empirical-probability matrix $\hat{\mathbf{P}}[N+1]$ incorporating $q_k[N+1]$, the quantized representation of the $(N+1)$st profile observation x_k, can be specified as

$$\hat{P}_{k,m}[N+1] = \begin{cases} \frac{N}{N+1}\hat{P}_{k,m}[N] + \frac{1}{N+1} & \text{if} \quad q_k[N+1] = m, \\ \frac{N}{N+1}\hat{P}_{k,m}[N] & \text{if} \quad q_k[N+1] \neq m. \end{cases} \tag{3}$$

Incremental updates of the empirical-probability matrix are thus trivial.

3.3.2. MPM Template: Sample-Penalty Component

Although the $K \times N_q$ matrix $\hat{\mathbf{P}}$ provides a compact, convenient statistical representation of the profiles in an MPM fingerprint, it does not provide a direct means for comparing profile observations to templates, or templates to templates. The second component of an MPM template is a $K \times N_q$ matrix $\mathbf{t}$ of sample penalties $t_{k,m}$ expressing the penalty to be assigned to any quantile observation at any sample, and enabling the calculation of scalar profile-to-template or template-to-template match scores as a sum of sample penalty values. In particular, the matrix $\mathbf{t}$ serves as a lookup table enabling the extremely efficient calculation of a profile-to-template match score Z as

$$Z = \frac{1}{\sqrt{K}} \sum_{k=1}^{K} t_{k,q_k}, \tag{4}$$

where q_k is the quantized profile to be scored against the template with sample-penalty matrix $t_{k,m}$, and where the motivation for scaling by $1/\sqrt{K}$ will be made clear shortly. Similarly, the matrix of $\mathbf{t}$ enables the efficient calculation of a pre-to-post-template match score S as a weighted sum of sample-to-sample penalties:

$$S = \frac{1}{\sqrt{K}} \sum_{k=1}^{K} \sum_{m=1}^{M} \hat{P}_{k,m} \cdot t_{k,m}, \tag{5}$$

where the collection of $\hat{P}_{k,m}$ empirical-probability values from the post-ambiguity template serve as weights on the corresponding $t_{k,m}$ sample-penalty values from the pre-ambiguity template. (Note that the template-to-template match score of (5) is exactly equivalent to the average of the profile-to-template match scores in (4) over the component post-ambiguity profiles used to calculate $\hat{P}_{k,m}$, against the pre-ambiguity template $t_{k,m}$.)

Given (4) and (5), construction of the sample-penalty component of an MPM template entails specification of a mapping from $\hat{P}_{k,m}$ to $t_{k,m}$. The mapping chosen for MPM is motivated by the match robustness premise in Section 2. In particular, the chosen mapping imposes limits on the contribution of any one sample to the overall match score in (4) or (5). It utilizes a matrix of "hedged" empirical probability estimates, constructed as follows:

$$\tilde{P}_{k,m} = \frac{N\hat{P}_{k,m} + \nu}{N + N_q\nu}, \tag{6}$$

where N is the number of profiles that were used to calculate $\hat{\mathbf{P}}$, and where ν is a free parameter that may be set to any positive value. Note that a positive ν prevents any $\tilde{P}_{k,m}$ from achieving a value of 0 or 1. Intuitively, (6) can be interpreted as a safeguard against drawing overly strong conclusions from a limited amount of data. Mathematically, (6) can be interpreted as the Bayesian combination of the observed distribution $\hat{\mathbf{P}}$ with a uniform prior to yield a posterior distribution $\tilde{P}_{k,m}$. In this interpretation, the parameter ν specifies the relative weight of the prior, relative to the observational term.

The $t_{k,m}$ are simply taken as a quadratic penalty $(1 - \hat{P}_{k,m})^2$, normalized by statistics calculated from $\tilde{P}_{k,m}$:

$$t_{k,m} = \frac{(1 - \hat{P}_{k,m})^2 - \hat{\mu}_k}{\hat{\sigma}_k}, \tag{7}$$

where

$$\hat{\mu}_k = \sum_{m=1}^{N_q} \tilde{P}_{k,m}(1 - \hat{P}_{k,m})^2 \tag{8}$$

and

$$\hat{\sigma}_k^2 = \sum_{m=1}^{N_q} \tilde{P}_{k,m}(1 - \hat{P}_{k,m})^4 - \hat{\mu}_k^2. \tag{9}$$

Note that (7) maps lower-probability events to higher penalties. Note also that $\hat{\mu}_k$ and $\hat{\sigma}_k$ are simply estimates of the mean and standard deviation of the quadratic penalty term, calculated using the hedged probability estimates $\tilde{P}_{k,m}$.

The implication of (7) is that in-class comparisons—that is, tests of individual profiles or templates against templates of the same class—will yield sample penalties with means of approximately 0 and standard deviations of approximately 1. Out-of-class comparisons—that is, tests of individual profiles or templates against templates of a different class—will yield sample penalties with unknown statistics, but which generally have significantly positive means. Assuming conditional independence of profile samples, the central limit theorem[6] can be used to show that in-class profile-to-template match scores Z calculated as in (4) and in-class template-to-template match scores S calculated as in (5) will be approximately standard normal. As with the sample penalties, no specific claim can be made for the statistics of out-of-class Z and S—other than the fact that they will almost certainly have significantly positive means. In other words, it is extremely likely that for any collection of profiles for a particular tracklet in a particular aspect bin, the in-class and out-of-class match-score distributions will be separable.

4. MPM SCORING DETAILS AND IMPLEMENTATION

The MPM templates generated by the target fingerprinter according to the process described in the previous section are used to score particular target-association hypotheses when requested by the KT. In general, the KT may request scoring for all joint or pairwise tracklet-to-tracklet association possibilities for a kinematic ambiguity

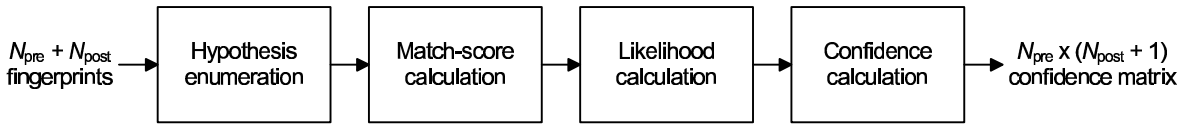

Figure 4. Block diagram of MPM scoring process.

in which N_{pre} tracklets enter a kinematically ambiguous state and N_{post} tracklets emerge. Additionally,the KT will generally require not only raw match scores, but also a measure of the likelihood or confidence associated with each possible joint or pairwise tracklet association.

The MPM scoring process is designed to provide match scores, likelihoods, and confidences for a general N_{pre}-in, N_{post}-out assignment problem. It consists of four fundamental steps, as illustrated in the block diagram of Figure 4. The first step is the enumeration of all possible joint tracklet-association hypotheses for the particular assignment problem at hand. The next step is the calculation of tracklet-to-tracklet match scores for each pairwise pre-to-post-tracklet assignment. The third step is the conversion of match scores to likelihoods, and the final step is the conversion of likelihoods into a matrix of confidences that can be reported to the KT. These steps are described in Sections 4.1 through 4.4, respectively. Note that although in general the KT will make its own tracklet-stitching decisions using the confidence information provided by the fingerprinter, it is possible to specify a simple threshold-based decision rule that will yield a target-identity declaration directly out of the tracker; this is discussed in Section 4.5.

4.1. Hypothesis Enumeration

The first step in the MPM scoring process is the enumeration of all joint pre-to-post-tracklet assignment hypotheses for the particular N_{pre}-in, N_{post}-out problem at hand. This calculation, while not the most conceptually or computationally challenging stage of the MPM scoring process, is still non-trivial. It is complicated by the need to account for "hiding targets," *i.e.*, targets that enter a kinematic ambiguity but do not emerge, or emerge from a kinematic ambiguity without entering it. (Although the presence of hiding targets is a necessity in assignment problems in which $N_{\text{pre}} \neq N_{\text{post}}$, it is also a possibility in problems in which $N_{\text{pre}} = N_{\text{post}}$, and must be considered in either case.) An additional complication is imposed by the combinatorial nature of the enumeration: although the number of possible *pairwise* pre-to-post tracklet assignments is linear in both N_{pre} and N_{post}, the number of possible *joint* pre-to-post tracklet assignments is combinatorial in both N_{pre} and N_{post}.

As an example of the accounting required for hypothesis enumeration, consider a two-in, three-out scenario. There are thirteen possible joint tracklet-association hypotheses: six in which both pre-ambiguity tracklets associate with post-ambiguity tracklets (leaving one post-ambiguity tracklet as a hiding target); six in which one of the pre-ambiguity tracklets associates with one of the post-ambiguity tracklets (leaving one pre-ambiguity tracklet and two post-ambiguity tracklets as hiding targets), and one in which no post-ambiguity tracklets associate with any pre-ambiguity tracklets (leaving both pre-ambiguity tracklets and all three post-ambiguity tracklets as hiding targets). Note that specification of all joint tracklet-association hypotheses becomes considerably more involved for larger N_{pre} or N_{post}.

4.2. Match Score Calculation

After hypothesis enumeration, the next step in MPM scoring is the calculation of match scores for each of the $N_{\text{pre}} \cdot N_{\text{post}}$ possible pairwise pre-to-post tracklet assignments. This is accomplished by calculating a collection of template-to-template match scores S according to (5) for each aspect-bin-to-aspect-bin comparison in each possible pre-to-post pairwise tracklet association. Note that because, in general, any tracklet will include profile observations from multiple aspect bins, this entails the calculation of a collection of scores S. In particular, if profiles are separated into N_ϕ aspect bins within each tracklet, then a single tracklet-to-tracklet match-score calculation may involve the computation of up to N_ϕ^2 template-to-template match scores. In the general case, in which profiles have been observed within only a subset of all available aspect bins, fewer than N_ϕ^2 template-to-template match scores will need to be generated.

Rather than combine template-to-template match scores to yield an overall scalar tracklet-to-tracklet match score, the full collection of all available template-to-template match scores is preserved as-is for further manipulation. In particular, the template-to-template match scores are combined in the likelihood stage, as described in the following section.

4.3. Likelihood Calculation

The likelihood calculation step of MPM scoring involves converting collections of template-to-template match scores for each pairwise tracklet comparison into likelihoods using estimates of the in-class and out-of-class template-to-template match-score distributions. If all in-class and out-of-class distributions were considered independent, this would suggest that we need to estimate $(N_{\text{pre}} + N_{\text{post}})^2 N_\phi^2$ separate distributions to capture the behavior of match scores for pairwise associations that do not involve hiding targets, and $(N_{\text{pre}} + N_{\text{post}})N_\phi^2$ additional distributions to capture the behavior of match scores for hiding-target associations. Furthermore, many of these distributions would correspond to tracklet or aspect pairings for which no observational data would be available from which to estimate distribution parameters.

We employ several simplifications to make this distribution estimation problem tractable. Simplification of the in-class distribution estimation is based on two convenient properties of in-class match scores. First, recall that in-class comparisons at the same aspect bin yield template-to-template match scores S that are approximately standard normal, by design. Second, empirical observation reveals that match scores obtained for cross-aspect comparisons of any particular target are approximately Gaussian, with means and standard deviations that increase with the the aspect-bin separation $\Delta\theta$. Based on these two observations, we model all in-class distributions as Gaussians and impose a parametric structure on their statistics. In particular, the means and standard deviations of all in-class distributions are specified as parametric functions incorporating a pre-specified prior that varies with $\Delta\theta$; the free parameters in the parametric representation can be estimated in closed form from the available data within each class. The out-of-class distributions for each target are also modeled as Gaussians, with statistics obtained by fitting available out-of-class data to a different parametric function representation. The result of the distribution estimation is thus a collection of $(N_{\text{pre}} + N_{\text{post}})N_\phi^2$ in-class mean and standard deviation estimate pairs and $(N_{\text{pre}} + N_{\text{post}})N_\phi^2$ out-of-class mean and standard deviation estimate pairs. We denote the in-class mean and standard deviation estimates as $\hat{\mu}_{i,i,j,k}$ and $\hat{\sigma}_{i,i,j,k}$, with i specifying the tracklet index $1, \ldots, (N_{\text{pre}} + N_{\text{post}})$, and with j and k specifying aspect-bin indices $1, \ldots, N_\phi$. We denote the out-of-class mean and standard deviation estimates as $\hat{\mu}_{i,0,j,k}$ and $\hat{\sigma}_{i,0,j,k}$.

Once the full collections of in-class and out-of-class statistic estimates are available, the likelihood of any particular template-to-template match score can be computed under any particular pairwise tracklet-association hypothesis. For instance, suppose we observe a match score $S_{i,j,k,l}$ for the comparison of the pre-ambiguity tracklet-i, aspect-bin-k matrix $\mathbf{t}$ with the post-ambiguity tracklet-j, aspect-bin-l matrix $\hat{\mathbf{P}}$. The likelihood of observing $S_{i,j,k,l}$ under a hypothesis specifying a pairwise association between pre-ambiguity tracklet i and post-ambiguity tracklet j is

$$L_{i,j,k,l} = \frac{1}{\hat{\sigma}_{i,i,k,l}\sqrt{2\pi}} \exp\left(-\frac{1}{2\hat{\sigma}_{i,i,k,l}^2}(S_{i,j,k,l} - \hat{\mu}_{i,i,k,l})^2\right). \tag{10}$$

Similarly, the likelihood of observing $S_{i,j,k,l}$ under a hypothesis indicating that pre-ambiguity tracklet i and post-ambiguity tracklet j are *not* associated is

$$L_{i,0,k,l} = \frac{1}{\hat{\sigma}_{i,0,k,l}\sqrt{2\pi}} \exp\left(-\frac{1}{2\hat{\sigma}_{i,0,k,l}^2}(S_{i,j,k,l} - \hat{\mu}_{i,0,k,l})^2\right). \tag{11}$$

In-class and out-of-class likelihoods of the match scores for each aspect-bin-to-aspect-bin comparison in each possible pre-to-post pairwise tracklet association (calculated in the previous MPM scoring step) are calculated using (10) and (11). Note that, in general, although they represent dual pre-to-post and post-to-pre representations of the same pairings, $S_{i,j,k,l} \neq S_{j,i,l,k}$ and $L_{i,j,k,l} \neq L_{j,i,l,k}$. Because both terms offer information, both are computed in the likelihood calculation step of MPM scoring.

4.4. Confidence Calculation

The calculated match-score likelihoods can be compared to obtain information about the relative certainties of the possible pairwise or joint tracklet-association hypotheses. For instance, consider the likelihood ratio

$$R_{i,j,k,l} = \frac{L_{i,j,k,l}}{L_{i,0,k,l}} \tag{12}$$

formed from the in-class and out-of-class likelihoods specified by (10) and (11). If $R_{i,j,k,l} > 1$, it is an indication that tracklet i and tracklet j share the same target identity, and should accordingly be associated. Alternatively, if $R_{i,0,k,l} < 1$, it is an indication that tracklet i and j have different identities, and should not be associated with each other.

The confidence calculation step of MPM scoring entails calculating $R_{i,j,k,l}$ for all in-class and out-of-class likelihood pairs calculated in the previous MPM scoring stage. As with the likelihoods, both the pre-to-post $R_{i,j,k,l}$ ratios and the post-to-pre $R_{j,i,l,k}$ ratios are calculated. MPM confidence assignment then involves combining ratios for each of the joint tracklet-association hypotheses enumerated for the particular N_{pre}-to-N_{post} assignment problem. This is first done by combining all aspect-bin pairings to yield pairwise whole-tracklet ratios of the form

$$R_{i,j} = \prod_{k,l} \alpha_{k,l} R_{i,j,k,l}, \tag{13}$$

where the $\alpha_{k,l}$ terms are pre-selected coefficients that can be chosen to weight comparisons at different aspects differently, or can simply be set to unity to treat all comparisons equally. The pairwise whole-tracklet ratio terms $R_{i,j}$ are in turn combined according to the pairwise components of each joint tracklet-association hypothesis to yield joint tracklet-association-hypothesis ratio terms. For instance, consider an assignment problem in which pre-ambiguity tracklets 1 and 2 enter a kinematic ambiguity and post-ambiguity tracklets 3, 4, and 5 emerge. For the joint tracklet-association hypothesis $H_{1:3,2:5,4:0}$ assigning tracklet 1 to tracklet 3, tracklet 2 to tracklet 5, and tracklet 4 to a hiding target, the joint tracklet-association ratio term is:

$$R(H_{1:3,2:5,4:0}) = \frac{R_{1,3}R_{2,5}}{R_{1,4}R_{1,5}R_{2,3}R_{2,4}} \cdot \frac{R_{3,1}R_{5,2}}{R_{3,2}R_{4,1}R_{4,2}R_{5,1}}. \tag{14}$$

This is simply a product of two compound terms, one representing the pre-to-post assignment and the second representing its post-to-pre dual. Each compound term is simply a fraction in which the numerator contains a term for every pairwise tracklet association that is made in the particular joint tracklet-association hypothesis, and the denominator contains a term for every pairwise tracklet association that is *not* made in the particular joint tracklet-association hypothesis. Note that if the given joint tracklet-association hypothesis $H_{1:3,2:5,4:0}$ is correct, then all of the numerator terms will tend to be greater than 1 and all of the denominator terms will tend to be less than 1, making the overall product significantly greater than 1. If the given joint tracklet-association hypothesis H is incorrect, then some of the numerator terms will tend to be less than 1 and some of the denominator terms will tend to be greater than 1. This will tend to reduce the overall product.

The complete set of N_H separate $R(H)$ terms, each calculated for a particular joint tracklet-association hypothesis according to the example of (14), are used to assign confidences to each joint tracklet-association hypothesis. In particular, for each of the N_H joint tracklet-association hypotheses H_i, we compute a confidence according to the formula

$$C(H_i) = \frac{R(H_i)}{\sum_{j=1}^{N_H} R(H_j)}. \tag{15}$$

This assigns a value between 0 and 1 to each of the N_H joint tracklet-association hypotheses, with higher values indicating more likely hypotheses. Note also that, by design, the full set of N_H joint-hypothesis confidences sums to 1. (The confidence metric described here has other useful properties, as well, as described elsewhere.[7])

If desired, the raw joint tracklet-association hypothesis confidences can be reported directly to the KT. Often, however, the KT requires specification of pairwise tracklet-association confidences. These can be directly obtained from the set of N_H joint tracklet-association confidences: the confidence in the pairwise association of pre-ambiguity tracklet i with post-ambiguity tracklet j can be specified as a function (typically the sum or maximum) of the $C(H_k)$ for which the corresponding joint tracklet-association hypothesis H_k involves the pairwise association of tracklet i and tracklet j. This process can be used to yield an $N_{\text{pre}} \times (N_{\text{post}} + 1)$ matrix $\mathbf{C}$ of pairwise confidences, in which the row-i, column-j entry corresponds to the confidence of associating pre-ambiguity tracklet i with post-ambiguity tracklet j if $j = 1, \ldots, N_{\text{post}}$, or to the confidence of assigning a "hiding target" label to pre-ambiguity tracklet i if $j = (N_{\text{post}} + 1)$. (If desired, this same process can be used to generate a dual $(N_{\text{pre}} + 1) \times N_{\text{post}}$ matrix of pairwise confidences keyed to the post-ambiguity tracklets instead of the pre-ambiguity tracklets.) The function combining the joint tracklet-association confidence terms $C(H_k)$ is typically chosen to ensure that each row of $\mathbf{C}$ sums to 1, so that the total confidence in all possible pairwise associations involving any pre-ambiguity tracklet is always unity.

4.5. Optional Decision Rule

If the target fingerprinter is being used in conjunction with a KT, then the fingerprinter would typically report the confidence matrix $\mathbf{C}$ to the KT and let the KT make the final tracklet-association declarations or deferrals. If the fingerprinter is being used as a standalone application, however, then the confidence matrix $\mathbf{C}$ can be used to produce target identity declarations. There are many ways in which this can be done. A simple but effective

decision rule that can be applied to $\mathbf{C}$ to yield target identity decisions is the application of a pre-selected decision threshold η, as follows.

Suppose the fingerprinter produces a confidence matrix $\mathbf{C}$ for an N_{pre}-in, N_{post}-out association problem. Let $C_{\max}(i)$ denote the maximum value observed in the ith row of $\mathbf{C}$, and let $j_{\max}(i)$ be the column in which this maximum value is achieved. Then for each pre-ambiguity tracklet i, we can specify an identity declaration or deferral according to the following rule:

$$\begin{array}{ll} \text{declare ``tracklet } i \text{ assigns to tracklet } j_{\max}(i)\text{''} & \text{if} \quad j_{\max}(i) \leq N_{\text{post}} \text{ and } C_{\max}(i) \geq \eta; \\ \text{declare ``tracklet } i \text{ is a hiding target''} & \text{if} \quad j_{\max}(i) = N_{\text{post}} + 1 \text{ and } C_{\max}(i) \geq \eta; \\ \text{defer decision} & \text{if} \quad C_{\max}(i) < \eta. \end{array} \tag{16}$$

This rule results in a target assignment only when the confidence in that assignment exceeds the pre-selected threshold η. Note also that if no possible assignments for a particular pre-ambiguity tracklet i exceed the specified threshold, then no explicit decision is made. By judicious selection of the decision threshold η, then, it is possible to increase the conditional probability of correct assignment (P_{ca}) at the expense of a larger deferral rate (P_{def}). For instance, selecting an η value near 1 will tend to result in a large P_{def}, but because all declarations will correspond to high confidences, P_{ca} will also tend to be large. Selecting an η value near $1/(N_{\text{post}} + 1)$ will tend to result in a small P_{def}, but at the expense of a greater fraction of incorrect declarations, *i.e.*, a smaller P_{ca}.

Finally, note that the decision rule of (16) does not safeguard against the possibility that multiple pre-ambiguity tracklets might be assigned to the same post-ambiguity tracklet, which obviously represents a physical impossibility. It is easy to preclude such an eventuality with a slightly more sophisticated decision rule. In practice, for η that are not close to $1/(N_{\text{post}+1})$, such "double assignment" is rarely observed.

5. EXPERIMENTAL RESULTS

We tested MPM using HRR profiles generated from SAR imagery in the publicly available MSTAR (Moving and Stationary Target Acquisition and Recognition) data set.[8,9] The MSTAR data was collected by Sandia National Laboratories for the Defense Advanced Research Projects Agency (DARPA). The publicly available MSTAR data are HH-polarization X-band SAR image chips with range and cross-range resolution of 0.3 m. This data set includes ten targets: BTR-60 transport, 2s1 gun, BRDM-2 truck, T-62 tank, ZIL-131 truck, ZSU-23-4 gun, D-7 bulldozer, BMP-2 tank, BTR-70 transport, and T-72 tank. For this study, we used SAR data collected at depressions of 15° and 17° to form HRR profiles which were then used as inputs to MPM. The profile-generation process is described in Section 5.1; the tracklet-association experimental setup and results are presented in Section 5.2.

5.1. Profile Generation

HRR profiles were generated from SAR imagery using a multi-step process. Each SAR chip was first cropped in azimuth to isolate the columns (*i.e.*, azimuth bins) containing target returns. This was accomplished using an automated segmentation algorithm: the maximum value in each column was computed, and all columns whose maxima exceeded a threshold of twice the mean of those maximum values were identified as target columns.

Once the target columns were segmented from the surrounding clutter, a Fourier transform was applied across azimuth within each target row (*i.e.*, range bin). This effectively reverses the azimuth-compression step of the underlying SAR image formation process. The resulting data is then processed to remove the Taylor window applied during the original SAR image formation process (characterized by a sidelobe level of -35.0 dB and an $\bar{n}$ value of 4). The result of this processing is a matrix in which each column could be considered a single-pulse HRR profile. To yield a single high-signal-to-noise-ratio profile, this matrix was converted into magnitude and then averaged across azimuth.

5.2. Tracklet Association Results

The profiles generated from the MSTAR data were used to demonstrate the performance of the MPM algorithm in several experiments, each involving simulation of a particular kinematic ambiguity scenario. In each experiment discussed here, a randomly selected set of vehicles enters a kinematic ambiguity, and the same set of vehicles emerges from the kinematic ambiguity. As discussed previously, the goal of MPM is to correctly associate the "pre-ambiguity" targets with the "post-ambiguity" targets.

We performed two sets of experiments: one set involving two-in, two-out kinematic ambiguities, and set one involving three-in, three-out kinematic ambiguities. In each case, the pre- and post-ambiguity targets were randomly selected from the 10 available MSTAR targets. The HRR profiles used as input data for each target

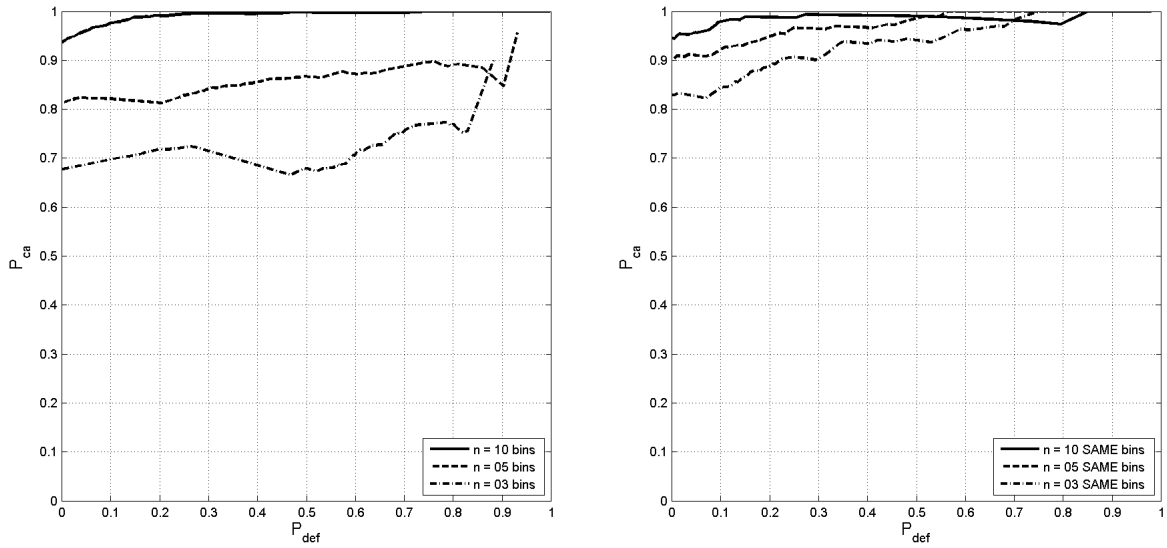

Figure 5. Two-in, two-out experimental results.

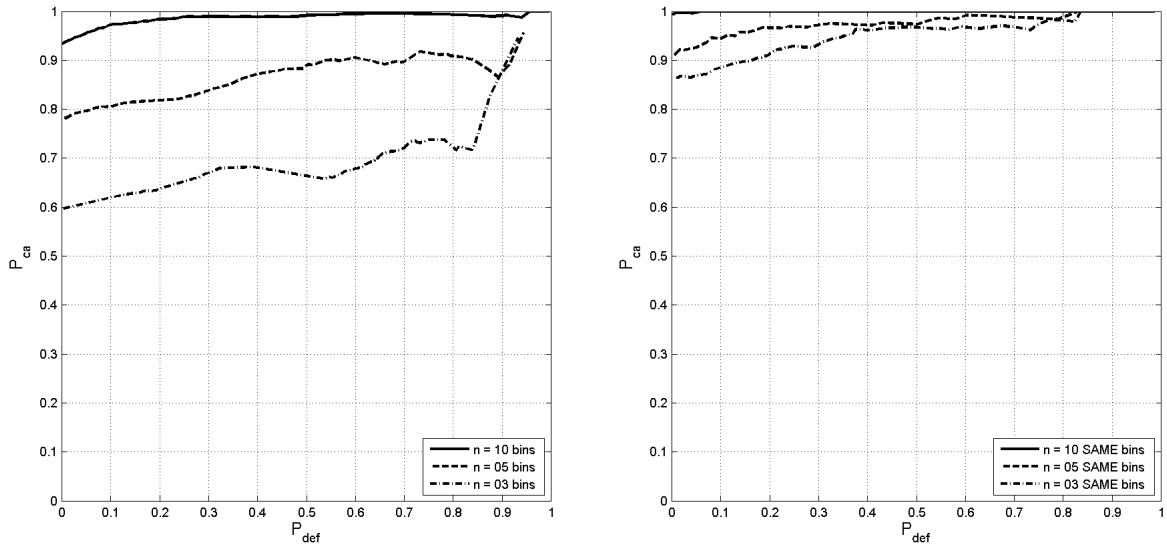

Figure 6. Three-in, three-out experimental results.

were distributed in n randomly selected 10° aspect bins, with six HRR profiles randomly selected per bin. For both the two-in, two-out and three-in, three-out scenarios, three separate values of n (3, 5, and 10) were tested. Additionally, to demonstrate the impact of correlation between pre-ambiguity aspect and post-ambiguity aspect, we performed one set of experiments in which the n randomly selected aspect bins were the same for all pre-ambiguity and post-ambiguity targets, and one set of experiments in which the n randomly selected aspect bins for the pre-ambiguity targets were chosen independently from the n randomly selected aspect bins for the post-ambiguity targets. In each trial, profile selection was done without replacement, to ensure that the same profile could not be present in both the pre- and post-ambiguity data.

The results of the full set of experiments are depicted in Figure 5 and Figure 6. The plots in Figure 5 depict the results of the two-in, two-out experiments, and the plots in Figure 6 depict the results of the three-in, three-out experiments. The left plot in each figure shows the results of the experiments in which the n aspect bins were chosen independently for pre- and post-ambiguity data, and the right plot shows the results of the experiments in which the n aspect bins were the same for pre- and post-ambiguity data. Each plot contains three traces, one for each of the three n values used (3, 5, and 10). Each plot is similar to a receiver operating characteristic (ROC) plot: the x-axis indicates probability of deferral (P_{def}), and the y-axis indicates declaration-conditional probability of correct assignment (P_{ca}); each point on a given trace is a particular P_{ca} and P_{def} pairing achieved using a particular confidence threshold η applied according to the decision rule specified in Section 4.5.

Several trends are immediately apparent from Figures 5 and 6. First of all, note that each trace illustrates the fundamental ROC-like trade-off between P_{ca} and P_{def}, in which a higher P_{ca} can be achieved at the expense of a higher P_{def}. (Although some of the traces exhibit local deviations from this behavior due to statistical variation and limited test data, the qualitative nature is apparent in the overall shape of each trace). Note also that in each plot, performance increases with more data—that is, as the number of aspect bins in which data is observed increases from 3 to 5, and from 5 to 10. Additionally, note that consistency in aspect between pre-ambiguity and post-ambiguity bins has a positive impact on performance, as is evident from a comparison of the left-hand and right-hand plots in each figure; this is attributable to the comparative ease of recognizing a target at a previously observed aspect, rather than at a randomly selected, and possibly as-yet-unobserved, aspect. Finally, note that performance is relatively consistent between the two figures: in other words, the inclusion of additional targets does not significantly impact MPM performance. Figures 5 and 6 illustrate that MPM is capable of excellent P_{ca} and P_{def} performance even with limited data, and in a variety of kinematic ambiguity scenarios.

6. SUMMARY

We have presented the multinomial pattern matching (MPM) algorithm for HRR target fingerprinting. The MPM target fingerprinting algorithm enables track stitching to facilitate long-term track maintenance of targets of interest even in the presence of kinematic ambiguities. MPM makes tracklet-association declarations based on characteristics inherent to each target class, rather than on the differences between classes. Its target class descriptions are based on information that is relatively stable within a target class—namely, the relative amplitudes of HRR-profile scattering responses, rather than the absolute amplitudes. MPM relies on a match metric that is robust to limited signature contamination and variability, ensuring that localized discrepancies between otherwise comparable target-class descriptions do not drive association decisions. MPM is structured and implemented to enable continuous real-time incremental training as new HRR data becomes available, and to enable tracklet-association hypotheses to be scored, and tracklet-association declarations to be made, at any point in time. MPM is not only capable of producing tracklet-association declarations, but also assigns a meaningful confidence value to all possible pairwise or joint tracklet-association hypothesis, enabling a more nuanced management of tracks and targets by a kinematic tracker, a sensor resource manager, or another application operating in conjunction with MPM. Performance was demonstrated on HRR profiles generated from the MSTAR data set. Results indicated excellent performance in several different scenarios.

ACKNOWLEDGMENTS

This work was sponsored in part by the Defense Advanced Research Projects Agency (DARPA) and was performed by Sandia National Laboratories. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy under contract DE-AC04-94AL85000.

REFERENCES

1. D. C. Schleher, *MTI and Pulsed Doppler Radar*, Artech House, Norwood, MA, 1991.
2. D. R. Wehner, *High Resolution Radar*, Artech House, Norwood, MA, 1987.
3. Y. Bar-Shalom and T. E. Fortmann, *Tracking and Data Association*, Academic Press, Orlando, FL, 1988.
4. K. M. Simonson, "Multinomial pattern matching: A robust algorithm for target identification," *Proceedings of the Automatic Target Recognizer Working Group*, 1997.
5. J. A. Richards, "Synthetic aperture radar," in *Encyclopedia of Optical Engineering*, R. G. Driggers, ed., pp. 2759–2771, Marcel Dekker Press, 2003.
6. W. Feller, *An Introduction to Probability Theory and its Applications*, vol. II, John Wiley & Sons, New York, second ed., 1971.
7. J. A. Richards, W. J. Bow, Jr., and B. K. Bray, "An informative confidence metric for ATR," in *Algorithms for Synthetic Aperture Radar Imagery X*, E. G. Zelnio and F. D. Garber, eds., *Proc. SPIE* **5095**, pp. 336–348, 2003.
8. "MSTAR (public) targets." Compact disc, 1998. Distributed by Air Force Research Laboratory, Dayton, OH.
9. "MSTAR/IU (public) mixed targets." Compact disc, 1998. Distributed by Air Force Research Laboratory, Dayton, OH.