Signature-Aided Tracking Using Association Hypotheses

Craig S. Agate and Kevin J. Sullivan

Toyon Research Corporation, 75 Aero Camino, Suite A, Goleta, CA, U.S.A

ABSTRACT

An algorithm is derived for signature-aided tracking which uses features (e.g. high-range resolution radar (HRRR) profiles), or *functions* of features, in addition to kinematic measurements to associate measurements to known tracks, clutter or new tracks. The approach taken here is to derive the probability of the measurement-to-track association hypotheses which incorporates the likelihood of features as well as the traditional approach of using the kinematic measurement likelihood. It is assumed that the probability density function (PDF) of the features (or some function of the features) is available from a library. The approach to probabilistically characterizing the PDF of the profiles relies on the availability of a class-specific library for each target type. The class-specific library of PDFs characterizes the profiles conditioned on the target class from which the profile originated and the aspect at which the profile was obtained. The algorithm is evaluated using the SLAMEM^{TM} simulation.

Keywords: association hypotheses, multi-hypothesis tracking, HRR profiles, features

1. INTRODUCTION

In the kinematic approach to the data association problem, an estimate of the target's state (kinematic values such as position, velocity, etc.), is utilized to predict the value of the next radar measurement of that target. Given a set of measurements from the radar, the measurements are assigned (associated) to the various tracks considering the *likelihood* that each measurement could have originated from each track. The likelihood is typically a kinematic measurement likelihood, meaning that the elements of the measurement vector are related to the target motion or position parameters and are independent of the *shape* of the target itself. Thus, the target is treated as a point target for the simple reason that traditional low-range resolution radar returns depend solely on the kinematic parameters of the target. With the advent of high-range resolution radar, "shape" information in addition to kinematic information on the target can be extracted from the measurements.

In many tracking scenarios, this newly available shape information, which we shall refer to as a signature or feature vector, is used to classify the target (i.e., identify (ID) the target). In this paper we propose a data association algorithm which utilizes the target signature in addition to the kinematic measurements to aid the association of measurements to tracks. The approach taken here parallels that of multi-hypothesis approaches (such as the N-scan algorithm,[1] the probabilistic data association (PDA) algorithm[2] or Joint probabilistic data association (JPDA) algorithm[3]) in the sense that hypotheses are formed assigning measurements to tracks. Then, the probability that each hypothesis is correct is calculated. The inclusion of a signature or feature likelihood in the hypothesis probability calculation is what marks the difference between the current approach and the traditional approaches. Depending on the various target types present, utilization of target signatures should significantly improve data association performance over traditional data association methods.

Implementing the SAT algorithm requires that a likelihood function of the features or signature be available. The derivation in subsequent sections is general enough to accommodate a likelihood function of the signature itself, a likelihood of *features* extracted from the signature,[4] or the likelihood of a *function* of the signature (i.e., a data statistic).

Further author information: (Send correspondence to C.S. Agate)
C.S. Agate: E-mail: cagate@toyon.com, Telephone: 1 805 968 6787 ext. 176
K.J. Sullivan: E-mail: ksullivan@toyon.com, Telephone: 1 805 968 6787 ext. 156

2. NOTATION

Before developing the SAT algorithm, we define some of the notation used and information required to implement the algorithm. The following information is assumed to be available before updating the tracks at time k.

- A set of known tracks:

$$\mathcal{T} = \{t_1, t_2, t_3, \ldots, t_N\} \tag{1}$$

  As we will assign measurements to known tracks, clutter or new tracks, we shall need a designation for clutter and new tracks in addition to the set of known tracks. Let t_0 represent a generic track to which measurements assumed to originate from either new tracks or clutter (also referred to as false alarms or spurious measurements) are assigned.

- A set of state estimates, $\hat{\mathbf{x}}$ *, and error covariance matrices, Σ_x, for each known track

$$\{\hat{\mathbf{x}}_1(k-1|k-1), \Sigma_{x_1}(k-1|k-1), \ldots, \hat{\mathbf{x}}_N(k-1|k-1), \Sigma_{x_N}(k-1|k-1)\} \tag{2}$$

  where

$$\hat{\mathbf{x}}_i(k-1|k-1) = E\{\mathbf{x}_i(k-1)|Z^{k-1}\} \tag{3}$$

  and Z^{k-1} represents the measurement history given by

$$Z^{k-1} = \{\mathbf{z}(1), \mathbf{z}(2), \ldots, \mathbf{z}(k-1)\} \tag{4}$$

- For each track there is a vector of class probabilities

$$P_{t_i}(k-1) = [P_{t_i}(\omega_1|\mathcal{E}^{k-1}) \ \cdots \ P_{t_i}(\omega_L|\mathcal{E}^{k-1})]^T \tag{5}$$

  where $P_{t_i}(\omega_l|\mathcal{E}^{k-1})$ represents the probability that track t_i belongs to an object from class ω_l given *all* the evidence up through time $k-1$. The evidence for our purposes will generally be the feature information although any information yielding clues about the identity of the tracks can be considered evidence.

- A *new* set of measurements consisting of both kinematic and feature information denoted

$$\begin{aligned} \mathbf{z}(k) &= \{\mathbf{z}_1(k), \mathbf{z}_2(k), \ldots, \mathbf{z}_M(k)\} \\ &= \{\mathbf{y}_1(k), \mathbf{s}_1(k), \mathbf{y}_2(k), \mathbf{s}_2(k), \ldots, \mathbf{y}_M(k), \mathbf{s}_M(k)\}, \end{aligned} \tag{6}$$

  where $\mathbf{y}_j(k)$ represents the j^{th} kinematic measurement and $\mathbf{s}_j(k)$ represents the corresponding feature measurement.

3. MEASUREMENT-TO-TRACK ASSOCATION ALGORITHM

Before deriving the association hypothesis probabilities, we first outline the basic steps of the SAT (signature-aided tracking) algorithm.

1. First, $\hat{\mathbf{x}}_i(k-1|k-1)$ and $\Sigma_{x_i}(k-1|k-1)$ for each track are utilized to implement a gating process resulting in a list, for each track, of measurements which *may* be associated with a particular track. Additionally, the feature likelihoods for associating the various detections to tracks is utilized as a gating metric. Note that a measurement may be associated with a known track, clutter or a new track. The result of this step is a *validation* matrix $\mathbf{\Omega}$ whose M columns represent measurements and whose $N+1$ rows represent tracks to which measurements may be associated. The elements of $\mathbf{\Omega}$ are determined by

$$\mathbf{\Omega}_{ij} = \begin{cases} 1 & \mathbf{z}_j(k) \text{ may have originated from } t_i \\ 0 & \mathbf{z}_j(k) \text{ low probability of originating from } t_i \end{cases} \tag{7}$$

*in the sequel, boldface letters such as $\mathbf{x}$ represent vectors

2. Based on the results of gating, form all possible measurement assocation hypotheses denoted $\Lambda_h(k)$ for $h = 1, 2, 3, \ldots, H$. An association hypothesis takes the set of measurements and decides the number which result from clutter or new tracks, n_U, and the number which result from known tracks, n_D. After hypothesizing the values of n_U and n_D, the actual mapping of measurements to tracks is made. This mapping can be summarized by a vector of length M (the number of measurements) whose elements contain the index of the track to which the measurement is assigned. An example is presented below to aid the understanding of this process. Assume that a set of 5 validated measurements is received, $\{\mathbf{z}_1(k), \mathbf{z}_2(k), \mathbf{z}_3(k), \mathbf{z}_4(k), \mathbf{z}_5(k)\}$, and that we have 3 known tracks. One possible hypothesis is to decide that of the 5 measurements, 2 belong to known tracks (i.e., $n_D = 2$) and 3 belong to either new tracks or are spurious (i.e., $n_U = 3$). Now that the number of clutter points, new tracks and detected tracks is chosen, a choice of which of the 5 measurements originate from clutter, which originate from new tracks and, finally, which measurements originate from *which* targets. The array [2 0 0 3 0] is one possible assignment given n_D and n_U and indicates that $\mathbf{z}_1(k)$ is associated with known track t_2, $\mathbf{z}_2(k)$, $\mathbf{z}_3(k)$ and $\mathbf{z}_5(k)$ are all associated with the generic track t_0, and $\mathbf{z}_4(k)$ is associated with known track t_3. Note that for this assignment, track t_1 is assumed undetected.
3. For each assocation hypothesis, $\Lambda_h(k)$, calculate the probability $P(\Lambda_h(k)|Z^k)$ for $h = 1, 2, \ldots, H$.
4. Calculate the probabilities of associating specific measurements with tracks as follows. Let $\theta_{t_i,j} = \{\text{Measurement } j \text{ originated from track } t_i\}$ where t_0 represents clutter or new tracks, and $t_1, t_2, \ldots, t_N$ represent known tracks. Find $P(\theta_{t_i,j}|Z^k)$ by summing the probabilities of *all* hypotheses in which measurement j is associated with track t_i. These calculated probabilities could be put into a matrix which is subsequently processed by a 2-D assignment algorithm[5] yielding a hard decision on how to associate the measurements to tracks. Alternatively, the hypothesis with the maximum probability of being correct may be chosen to update the track estimates.

4. DETERMINING THE HYPOTHESIS PROBABILITIES

We now derive expressions for determining the various hypothesis probabilities. The derivation here parallels that used in multihypothesis tracking save for the inclusion of feature measurements. We begin by separating the measurements $\mathbf{z}(k)$ into the kinematic and feature parts.

$$\begin{aligned} P(\Lambda_h(k)|Z^k) &= P(\Lambda_h(k)|S^k, Y^k) \\ &= P(\Lambda_h(k)|\mathbf{s}(k), \mathbf{y}(k), S^{k-1}, Y^{k-1}) \\ &= \frac{p(\mathbf{y}(k), \mathbf{s}(k)|\Lambda_h(k), Y^{k-1}, S^{k-1}) P(\Lambda_h(k)|Y^{k-1}, S^{k-1})}{p(\mathbf{y}(k), \mathbf{s}(k)|Y^{k-1}, S^{k-1})}, \end{aligned} \tag{8}$$

where $h = 1, 2, \ldots, H$. Note that the denominator in (8) is merely a normalizing factor which is independent of the particular hypothesis. Thus, we simplify (8) to

$$P(\Lambda_h(k)|Z^k) = \beta p(\mathbf{y}(k), \mathbf{s}(k)|\Lambda_h(k), Y^{k-1}, S^{k-1}) P(\Lambda_h(k)), \tag{9}$$

where we have eliminated the conditioning of $\Lambda_h(k)$ on Y^{k-1} and S^{k-1} since without conditioning on the current measurement as well, the truth of $\Lambda_h(k)$ is unaffected by Y^{k-1} or S^{k-1}. Next (9) is manipulated into

$$P(\Lambda_h(k)|Z^k) = \beta p(\mathbf{s}(k)|\mathbf{y}(k), \Lambda_h(k), Y^{k-1}, S^{k-1}) p(\mathbf{y}(k)|\Lambda_h(k), Y^{k-1}, S^{k-1}) P(\Lambda_h(k)). \tag{10}$$

At this point a discussion of the factors influencing the densities of both the profiles, $\mathbf{s}(k)$ and kinematic measurements, $\mathbf{y}(k)$, is necessary. The density of a kinematic measurement is typically assumed to be Gaussian distributed with a mean and covariance determined by the Kalman filter estimates of the track with which that measurement is associated. Thus, the density of $\mathbf{y}(k)$ depends on $\Lambda_h(k)$, which associates $\mathbf{y}(k)$ to a track. The density also depends on the history of kinematic measurements, Y^{k-1} which yields an estimate of the target dynamics. Without incorporating the orientation of the target into the dynamic state and using the profiles to estimate it, the history of feature measurements, S^{k-1} has no influence on the density of $\mathbf{y}(k)$. A feature measurement is a function of the target type and aspect angle with which the radar illuminates the target. The aspect angle, in turn, is a function of the relative position of the target and radar, which is known, and the orientation of the target which must be estimated from the kinematics. Thus, the density of $\mathbf{s}(k)$ depends on all kinematic measurements, Y^k, a hypothesis, $\Lambda_h(k)$, which associates the feature with a specific track and, finally, depends on S^{k-1} which influences the target type as quantified by the vector of class probabilities, $P_i(k-1)$.

Before continuing with the derivation of the different densities in (10), a comment on an alternative application of conditional probability to (9) is in order. We could have applied the rules of conditional probability to (9) to yield

$$P(\Lambda_h(k)|Z^k) = \beta p(\mathbf{y}(k)|\mathbf{s}(k), \Lambda_h(k), Y^{k-1}, S^{k-1}) p(\mathbf{s}(k)|\Lambda_h(k), Y^{k-1}, S^{k-1}) P(\Lambda_h(k)). \tag{11}$$

While this is a perfectly valid application of conditional probability rules, the resulting densities are less suitable for our purposes. The density $p(\mathbf{s}(k)|\Lambda_h(k), Y^{k-1}, S^{k-1})$ lacks conditioning on the most recent kinematic measurement. Thus, the orientation estimate at the current time (k) (necessary to look up the feature PDF in the library) must be estimated from past data unlike the feature density in (10) which includes conditioning on the most recent measurement. The density $p(\mathbf{y}(k)|\mathbf{s}(k), Y^{k-1}, S^{k-1}, \Lambda_h(k))$ poses other difficulties. Note that it is conditioned on the current feature measurement $\mathbf{s}(k)$. While $\mathbf{s}(k)$ contains information about the orientation of the target, which in turn affects the density of the kinematic measurement, accessing this information requires a search over a library of profiles indexed by target type and orientation with respect to the radar. Clearly, this process is not amenable to real-time applications.

While the kinematic and feature measurements have been discussed without reference to the specific type of measurement (e.g., synthetic aperture radar (SAR) profile or high-range resolution radar (HRRR) profile), the particulars associated with the various measurement types will be addressed as needed. We now continue with our derivation of the hypothesis probabilities by determining the individual probabilities and densities in (10).

4.1. Kinematic Measurement Density, $p(\mathbf{y}(k)|\Lambda_h(k), Y^{k-1})$

First note that $\mathbf{y}(k) = \{\mathbf{y}_1(k), \mathbf{y}_2(k), \ldots, \mathbf{y}_M(k)\}$ is the set of all current kinematic measurements. Conditioning the density on a specific track-to-measurement association, we shall assume that the measurements are independent. The density of $\mathbf{y}(k)$ is then written

$$p(\mathbf{y}(k)|\Lambda_h(k), Y^{k-1}) = \prod_{j=1}^{M} p(\mathbf{y}_j(k)|Y^{k-1}, \Lambda_h(k)). \tag{12}$$

The form of the individual densities in (12) depends on whether the measurement is associated with a known track, clutter or unknown track and also depends on the *type* of measurement received. For a measurement associated with an ground moving target indicator high-range resolution (HRRGMTI) radar, we assume that the kinematic part consists of a range, range rate and bearing measurement (i.e., $[R\ \dot{R}\ \theta]^T$), and that the signature part consists of an HRR profile. For a measurement associated with a SAR image, however, the kinematic part consists of position in Cartesian coordinates (i.e., $[x\ y\ z]$). Given a SAR image, we assume that the following processing steps take place

1. The 3-D position of an object detected in the image is extracted (local terrain maps are available).
2. The relative orientation of the object in the image is estimated.
3. A profile of the object in the image is extracted.

4.1.1. $\mathbf{y}_j(k)$ originates from a known track t_i for $1 \leq i \leq N$

In this case the measurement is assumed to be Gaussian distributed about the predicted measurement, $\hat{\mathbf{y}}_j(k|k-1) = E\{\mathbf{y}_j(k)|\Lambda_h(k), Y^{k-1}\}$, with a covariance matrix, $\Sigma_{y_j}(k|k-1)$. However, it is not a true Gaussian but a truncated Gaussian since the measurements were first validated by a gate. Accordingly, the density must be scaled by the probability that the measurement "falls" within the gate of some track (denoted P_G)

$$p(\mathbf{y}_j(k)|Y^{k-1}, \Lambda_h(k)) = \frac{1}{P_G(2\pi)^{\frac{q}{2}}|\Sigma_{y_j}(k|k-1)|^{\frac{1}{2}}} \exp(-\frac{1}{2}\nu_{t_i,j}(k)^T \Sigma_{y_j}^{-1}(k|k-1)\nu_{t_i,j}(k)), \tag{13}$$

where $\mathbf{y}_j(k) \in \mathbb{R}^q$ and $\nu_{t_i,j}(k) = \mathbf{y}_j(k) - \hat{\mathbf{y}}_j(k|k-1)$.

The evaluation of P_G is based on the type of validation used. While we often speak of the gating as spatial (i.e. depends on position in $\mathbb{R}^3$), it is really a gate on the measurement space of which $\mathbf{y}_j(k)$ is an element. The measurement space may include velocities as well as positions in $\mathbb{R}^3$. For example, the measurement vector for a GMTI radar is $\mathbf{y} = [R\ \dot{R}\ \theta]^T$ where R is the range to the target, θ is the bearing to the target and $\dot{R}$ is the range rate (i.e., radial velocity) of target. Thus, P_G may be calculated using

$$P_G = \int_{V_G} \cdots \int p(\mathbf{y}_j(k)|Y^{k-1}, \Lambda_h(k)) d\mathbf{y}_j(k) \tag{14}$$

where V_G is the region of the gate.

4.1.2. $\mathbf{y}_j(k)$ originates from clutter or a new track (track t_0)

The standard assumption is that clutter, or a measurement from a new track, is likely to appear anywhere, equally likely, within the validation gate. Thus, we say it is spatially distributed uniformly. Denoting the volume of the validation region as V_G, the PDF of a measurement originating from clutter is given by

$$p(\mathbf{y}_j(k)|Y^{k-1}, \Lambda_h(k)) = \frac{1}{V_G}. \tag{15}$$

Again, the term "volume" may mislead us into thinking that the validation region is a subset of 3-dimensional space. The validation region is simply a subset of the measurement space which, again, may contain velocity as an element of the space in addition to any physical region (2-dimensional in the case of ground targets, 3-dimensional in the case of airborne objects or 1-dimensional in the case of road-clamped ground targets).

4.2. Class-Dependent Feature Measurement Density

In the class-dependent approach to determining the feature likelihood, the PDF of the feature (or a statistic derived from it) is conditioned on the specific class type from which the feature is assumed to have originated, and the aspect from which the target was illuminated by the radar. Given a class type and aspect, then, the appropriate PDF is accessed from a library of PDFs.

As in the case of kinematic measurements, $\mathbf{s}(k) = \{\mathbf{s}_1(k), \mathbf{s}_2(k), \ldots, \mathbf{s}_M(k)\}$ is a set of feature measurements which are assumed independent when conditioned on a specific hypothesis, $\Lambda_h(k)$. Thus,

$$p(\mathbf{s}(k)|\Lambda_h(k), S^{k-1}, Y^k) = \prod_{j=1}^{M} p(\mathbf{s}_j(k)|\Lambda_h(k), S^{k-1}, Y^k). \tag{16}$$

Again, just as in the case of kinematic measurements, a feature measurement may be assigned to one of the N known tracks, or associated with track t_0 if it is assumed to have originated from a new track or simply be a spurious measurement. Before continuing, note that in the derivation that follows, the PDF of the feature vector is sought. However, we may map the feature vector into a statistic for which we have a PDF. An example would be the "electronic length" of an object which can be derived from the profile. Suppose that $\xi_j(k) = g(\mathbf{s}_j(k))$ is a statistic derived from the feature measurement. Alternatively, then, the likelihood for the vector of statistics, $\boldsymbol{\xi}(k)$, given by

$$p(\boldsymbol{\xi}(k)|\Lambda_h(k), S^{k-1}, Y^k) = \prod_{j=1}^{M} p(\xi_j(k)|\Lambda_h(k), S^{k-1}, Y^k), \tag{17}$$

is used in place of the likelihood of features in (16). It is likely that $\boldsymbol{\xi}(k)$ is *not* a sufficient statistic since information is generally lost in computing attributes like the electronic length. This loss of information must be accepted when the PDFs of the feature vectors themselves are not available.

4.2.1. $\mathbf{s}_j(k)$ originates from a known track t_i for $1 \leq i \leq N$

We now derive the PDF of a feature vector which is associated with a known track t_i. Assuming that the set of object classes is mutually exclusive and exhaustive (i.e. that a target *must* belong to (only) one of the L classes), we can write, using the total probability law,

$$\begin{aligned} p(\mathbf{s}_j(k)|\Lambda_h(k), Y^k, S^{k-1}) &= \sum_{l=1}^{L} p(\mathbf{s}_j(k), t_i \in \omega_l|\Lambda_h(k), Y^k, S^{k-1}) \qquad (18) \\ &= \sum_{l=1}^{L} p(\mathbf{s}_j(k)|t_i \in \omega_l, \Lambda_h(k), Y^k, S^{k-1}) P(t_i \in \omega_l|\Lambda_h(k), Y^k, S^{k-1}), \end{aligned}$$

where we have assumed that $\Lambda_h(k)$ associates $\mathbf{s}_j(k)$ with track t_i. By summing over the target classes, we have conditioned the feature density on the target type, which is one of the variables upon which the density directly depends. The other variable affecting the density is the aspect with which the radar "views" the target. Since the database of signature PDFs is assumed to be accessed by class type and aspect angle, we must further manipulate (18). Denote the aspect, which is specified by an elevation and an azimuth, of target t_i as $\boldsymbol{\psi}_{t_i}(k)$. Conditioning on the aspect angle leads to the following expression

$$p(\mathbf{s}_j(k)|\Lambda_h(k), Y^k, S^{k-1}) = \sum_{l=1}^{L} \left\{ \int \cdots \int p(\mathbf{s}_j(k)|t_i \in \omega_l, \Lambda_h(k), \boldsymbol{\psi}_{t_i}(k)) p(\boldsymbol{\psi}_{t_i}(k)|\Lambda_h(k), Y^k, S^{k-1}) d\boldsymbol{\psi}_{t_i}(k) \right\} P(t_i \in \omega_l|S^{k-1}) \tag{19}$$

Note that $P(t_i \in \omega_l|S^{k-1})$ $l = 1, 2, \ldots, L$ are simply the prior class probabilities for track t_i. The feature likelihood conditioned on class type and aspect angle, $p(\mathbf{s}_j(k)|t_i \in \omega_l, \Lambda_h(k), \boldsymbol{\psi}_{t_i}(k))$, is assumed available from a library of signature/feature PDFs. The remaining term, $p(\boldsymbol{\psi}_{t_i}(k)|\Lambda_h(k), Y^k, S^{k-1})$, is the PDF of the aspect angle(s) given the track measurement history (including the current kinematic measurement) from which the signature is assumed to have originated. The aspect angle can be estimated from knowledge of the track state, but there is uncertainty in the track state estimate as well as the assumed knowledge of the sensor platform's position and orientation. This uncertainty must be reflected in the PDF. Due to the number of different possible errors (sensor position and orientation errors, road network errors, etc.), only an approximation of $p(\boldsymbol{\psi}_{t_i}(k)|\Lambda_h(k), Y^k, S^{k-1})$ is possible. For example, a gaussian PDF centered about an estimate of $\boldsymbol{\psi}_{t_i}$ could be used in (19). The estimate of $\boldsymbol{\psi}_{t_i}$ can be derived from a combination of the track estimate and an estimate arrived at through road clamping.

4.2.2. $\mathbf{s}_j(k)$ originates from a new track or clutter

A received measurement (kinematic and feature) may originate from clutter (i.e., false alarm) or from a previously unknown target. It is assumed that processing of the radar returns eliminates any measurements whose signature is obviously not a target. Thus, given a set of "valid" measurements from the radar there is no way to distinguish a signature as having originated from a new track or clutter. The formula in (18) is still valid; however, since t_i (the track to which the feature is assigned) represents a possibly new track, there is no history of measurements. Eliminating the conditioning on past measurements in (18) yields the following formula for the feature likelihood

$$p(\mathbf{s}_j(k)|S^{k-1}, Y^{k-1}, \mathbf{y}(k), \Lambda_h(k)) = \sum_{l=1}^{L} p(\mathbf{s}_j(k)|t_0 \in \omega_l, \mathbf{y}(k)) P(t_0 \in \omega_l), \tag{20}$$

where we have let $t_i = t_0$ since the measurement is assigned to a previously unknown track. The term, $P(t_0 \in \omega_l)$, represents the probability that a track belongs to an object of class ω_l given no other knowledge. This can be assumed to be uniformly distributed over the set of possible target classes. Alternatively, using prior knowledge on the type of targets found in a particular region/scenario, a different distribution could be utilized. Again, the feature likelihood, conditioned on class type and aspect, is assumed to be available from a database. Thus, we condition the feature likelihood in (20) on the apsect angle to obtain the equivalent to (19) for the case of a feature originating from a previously unknown track.

$$\begin{aligned} & p(\mathbf{s}_j(k)|\Lambda_h(k), Y^k, S^{k-1}) = \\ & \qquad \sum_{l=1}^{L} \left\{ \int \cdots \int p(\mathbf{s}_j(k)|t_0 \in \omega_l, \Lambda_h(k), \boldsymbol{\psi}_{t_i}(k)) p(\boldsymbol{\psi}_{t_i}(k)|\mathbf{y}(k)) d\boldsymbol{\psi}_{t_0}(k) \right\} P(t_0 \in \omega_l) \end{aligned} \tag{21}$$

The azimuth and elevation angles can be assumed to be distributed uniformly without any other information. However, given $\mathbf{y}(k)$ we have *some* information about the object aspect angles. The elevation, albeit a noise corrupted version, can be found directly by the range and bearing measurements. While the azimuth is unknown, it is theoretically possible to find the distribution of the azimuth conditioned on knowing the range rate (i.e., the radial component of the target velocity $\mathbf{v}_R$) and the distribution of the vehicle speed. Mathematically, we wish to determine $p(\psi_{az}|\mathbf{v}_R)$ for which we need a distribution on the vehicle speed $\|\mathbf{v}\|$. For example, the vehicle speed might be assumed to have a beta distribution on $[0, \|\mathbf{v}\|_{max}]$ where $\|\mathbf{v}\|_{max}$ is the maximum vehicle speed. Using this density and the relationships between ψ_{az}, $\|\mathbf{v}_R\|$, $\|\mathbf{v}_T\|$ (the tangential velocity component) and $\|\mathbf{v}\|$, the distribution of ψ_{az} can be found.

4.3. Unconditional Hypothesis Probability

The remaining term for calculating the hypothesis probabilities in (8) is the unconditional probability $p(\Lambda_h(k))$. Carefully defining all the information implicit in "$\Lambda_h(k)$", facilitates the computation of this probability. Defining the following events

- $M = \{$there are M validated measurements$\}$
- $n_U = \{n_U$ of the M measurements originate from false or unknown tracks$\}$
- $n_D = \{n_D$ of the M measurements originate from detected, known tracks$\}$
- ζ is the mapping of measurements to specific tracks,

allows us to write a hypothesis as the set of events given by

$$\Lambda_h(k) \equiv \{M,\ n_U,\ n_D,\ \zeta\}. \tag{22}$$

Calculating the probability of a hypothesis is accomplished using conditional probability rules as

$$\begin{aligned} p(\Lambda_h(k)) &= p(\zeta, M, n_U, n_D) \\ &= p(\zeta|M, n_U, n_D) p(M, n_U, n_D) \\ &= p(\zeta|M, n_U, n_D) p(n_U, n_D|M) p(M) \\ &= p(\zeta|M, n_U, n_D) p(n_U|M) p(n_D|M) p(M), \end{aligned} \tag{23}$$

where $p(M)$ is the same for all hypotheses and gets absorbed into any normalization constant. Given specific values for n_D, M and n_U, let N_ζ represent the number of different mappings of measurements to the particular n_D detected tracks. Given the specific n_D tracks detected, there are $N_\zeta = \frac{M!}{(M-n_D)!}$ ways that the same set of targets can be detected given there are M total measurements. Assuming that these mappings are equally likely, $p(\zeta|M, n_U, n_D) = \frac{1}{N_\zeta} = \frac{(M-n_D)!}{M!}$. The second term in (23) is the probability of receiving n_U clutter points or new tracks which is often modelled as a Poisson random variable. Under this assumption, $p(n_U|M)$ is given by

$$p(n_U|M) = \frac{(\lambda_U V_G)^{n_U}}{n_U!} \exp(-\lambda_U V_G), \tag{24}$$

where n_U is the number of measurements assumed to have originated from new tracks or clutter (i.e., false alarms), λ_U is a density factor (the number of expected returns per volume) and V_G is the volume of the measurement gate. Finally, $p(n_D|M)$ is simply given by

$$p(n_D|M) = (P_D P_G)^{n_D}(1 - P_D P_G)^{N-n_D}, \tag{25}$$

where N is the number of known tracks. Putting all these formulas together results in the following *a priori* probability that hypothesis $\Lambda_h(k)$ is correct

$$p(\Lambda_h(k)) = \frac{(M-n_D)!}{M!} \left(\frac{(\lambda_U V_G)^{n_U}}{n_U!} \exp(-\lambda_U V_G) \right) (P_D P_G)^{n_D}(1 - P_D P_G)^{N-n_D}. \tag{26}$$

5. SIGNATURE LIKELIHOODS

We now discuss the determination of the signature PDFs. As mentioned previously, signature PDFs can be built for the signature itself. Thus, for the case of HRR profiles, the signature is a vector of values where each value represents the value of the profile in a particular range bin (see Figure 1). Consider each range bin value to be a random variable. If we assume that the noise in each range bin is independent, then we can build a PDF for the entire profile as the product of the PDFs for the individual elements of the profile. Thus, the PDF of the profile $\mathbf{s}(k)$ in Figure 1 is given by

$$p(\mathbf{s}(k)|\omega, \boldsymbol{\psi}(k)) = \prod_{j=1}^{m} p(s_j(k)|\omega, \psi_{az}, \psi_{el}), \tag{27}$$

where ψ_{az} and ψ_{el} are the elements of the aspect vector corresponding to the azimuth and elevation, respectively. A parametric form for the PDF of the range bin values is assumed (e.g., Gaussian, Beta) and the parameters determined using a training set of profiles for each class and aspect.

An alternative to finding the PDF of the signature itself is to find the PDF of a *function* of the signature. The function represents a data statistic which, for most cases will *not* be a sufficient statistic. However, as long as likelihood functions can be determined for the data statistic for each class and aspect, we can utilize this as the signature likelihood function within the framework for the feature-aided tracker. As an example of using a statistic in place of the feature, suppose that given the profiles of two targets from the same class and viewed from the same aspect, we compute some score which measures the "similarity" of the profiles. Due to statistical variations in features within one class, the resulting score must be statistically modelled by a PDF. Calculating the feature likelihood for our association problem, then, could be done in the following manner. Given a class type and aspect estimate, the score between $\mathbf{s}_j(k)$ and a "typical" profile from the same class and viewed from the same aspect is computed. The resulting score is substituted into the PDF, thereby generating the likelihood $p(\xi_j(k)|\Lambda_h(k), S^{k-1}, Y^k)$. This approach was implemented in the algorithm testbed (see Section 6). The statistic used was the SHARP MSE score.[6]

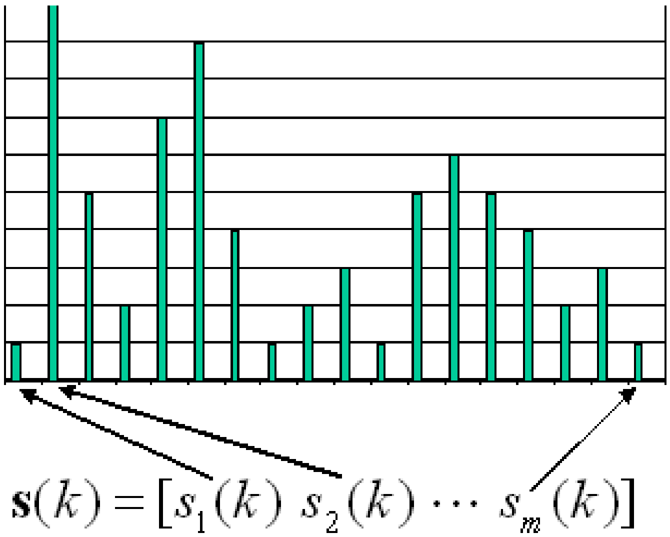

Figure 1: HRR profile as a vector of random variables.

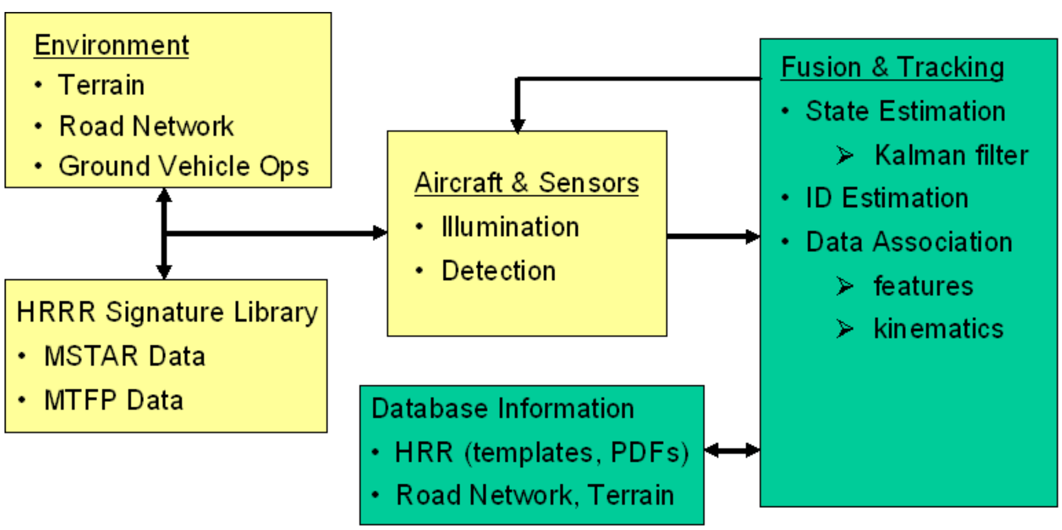

Figure 2: Block diagram of simulation environment.

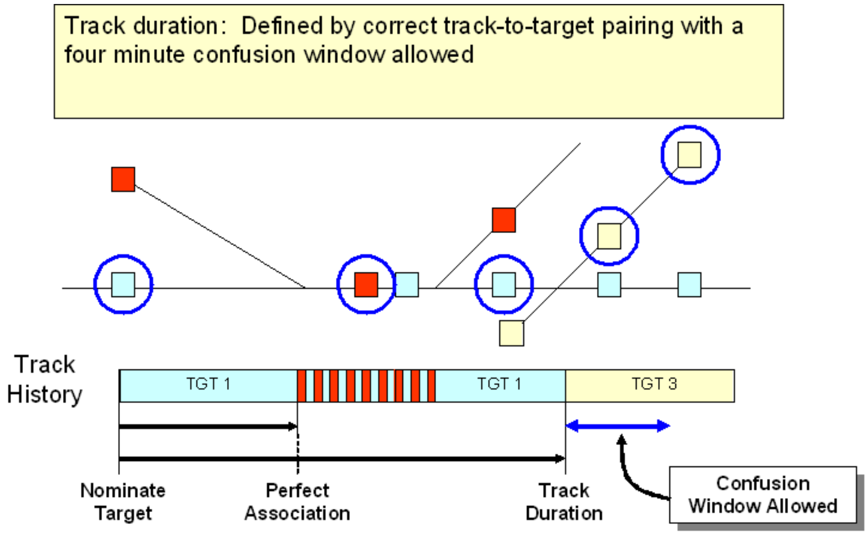

Figure 3: Illustration of performance metrics used to measure track time.

6. SIMULATION RESULTS

The algorithm was tested in a simulation environment called SLAMEM^{TM} which incorporates motion models for ground vehicle operations, sensor models and road networks. This tool has been used for analysis and as an algorithm testbed. A diagram of the environment is shown in Figure 2. In addition to ground vehicle operations, road networks, terrain and sensor models, a database of HRRR measurements has been incorporated into the simulation. To test the algorithm, two HRRGMTI radars at an altitude of $66,000$ feet pass frames of detections (10 second revisit rate) to the fusion and tracking algorithm which updates estimates of each track's kinematic state as well as class type. In one scenario, there were two target vehicles belonging to the same class and a varying number of background vehicles all belonging to the same class which was distinct from the target vehicles' class. The vehicles move about on 180 km of road in an area $12x12$ km^2.

The background vehicles move around continuously for 1.25 hours while the two target vehicles begin moving at a time uniformly distribed over the first 15 minutes of the simulation. The performance measure is based on the track duration of the target vehicles. Two measures of track duration were used. A *perfect track time* is defined as the time measured from the track initialization of a target vehicle to the moment that a detection from some *other* vehicle is used to update the track. In other words, the perfect time is a measure of the track *purity*. The other measure of track duration that was used allows for track impurity in the following way. The *metric track duration* represents the time from the track initialization of a target vehicle to the point at which the track has *not* been updated by a target detection in 4 minutes *minus* a 4 minute window. Thus, for the metric track duration, detections from other vehicles are allowed to update the track of the target vehicle as long as detections from the target vehicle are subsequently used to update the target vehicle's track within 4 minutes of being updated by the detection of another vehicle. An illustration of the two different track duration measures is shown in Figure 3.

A comparison of the metric track times for the feature-aided tracker with the kinematic tracker is shown in Figure 4. The number of background vehicles is varied and the metric track duration obtained by averaging over Monte Carlo runs. Clearly, for this particular tracking algorithm, features significantly improved the tracker.

To compare the difference between the perfect track time and the metric track time, the times are plotted in Figure 5 for the case of 60 background vehicles. The plot is the probability distribution function of the track duration and is plotted both for the feature-aided tracker and pure kinematic tracker. It is interesting to note that for the kinematic tracker, there is little difference between the perfect track time and metric track time. This suggests that in the case of pure kinematic tracking, once detections from other vehicles are incorporated

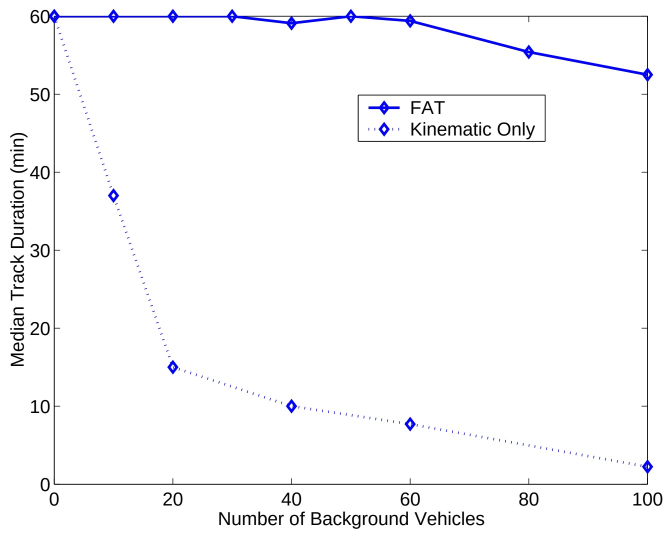

Figure 4: Metric track duration of feature-aided tracking versus pure kinematic tracking

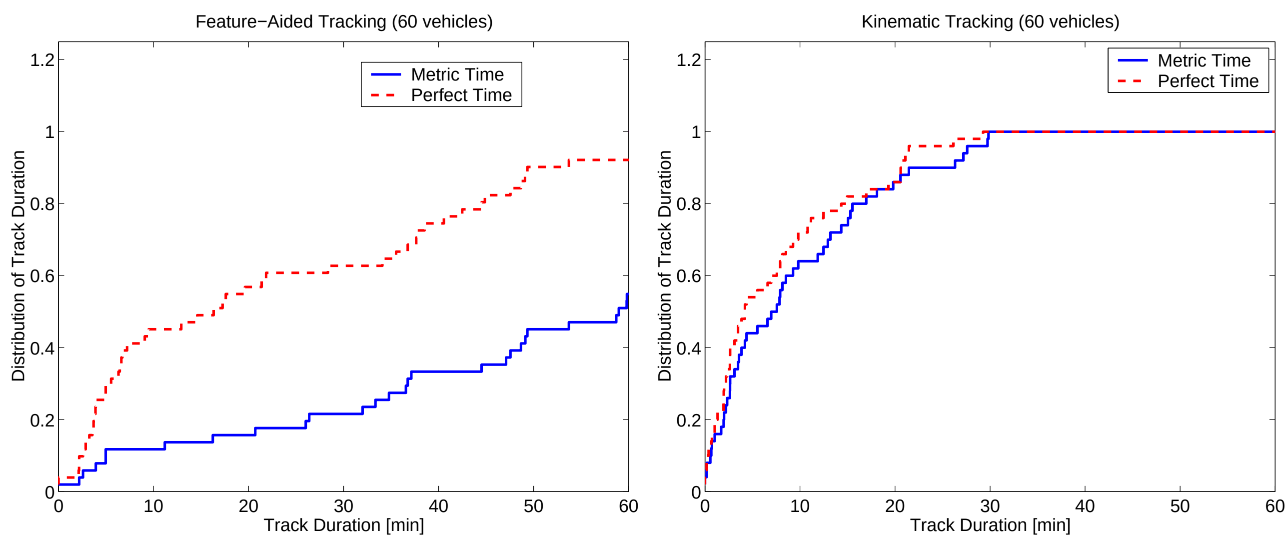

Figure 5: Comparison of perfect track times and metric track times for 60 vehicles

into the track of a target vehicle, the track is most likely going to be lost; whereas, in the case of feature-aided tracking, the track can often recover and continue tracking the correct vehicle.

ACKNOWLEDGMENTS

The authors wish to thank Alan Wood (AFRL/SN) for his technical input and sponsorship of this research. The authors also wish to thank Dr. Harry Burger for the development of the kinematic tracker.

REFERENCES

1. R. Singer, R. Sea, and K. Housewright, "Derivation and evaluation of improved tracking filters for use in dense multitarget environments," *IEEE Transactions on Information Theory* **IT-20**, pp. 423–432, July 1974.

2. T. Fortmann, Y. Bar-Shalom, and M. Scheffe, "Sonar tracking of multiple targets using joint probabilistic data association," *IEEE Journal of Oceanic Engineering* **OE-8**, pp. 173–184, July 1983.
3. Y. Bar-Shalom and T. Fortmann, *Tracking and Data Association*, Academic Press, Boston, MA, 1988.
4. R. Mitchell and J. Westerkamp, "Robust statistical feature based aircraft identification," *IEEE Transactions on Aerospace and Electronic Systems* **AES-35**, pp. 1077–1093, July 1999.
5. K. Pattipati, R. Popp, and T. Kirubarajan, "Survey of assignment techniques for multitarget tracking," in *Multitarget-Multisensor Tracking: Applications and Advances Volume III*, Y. Bar-Shalom and W. Blair, eds., pp. 77–160, Artech House, Norwood, MA, 2000.
6. M. Ressler, R. Williams, D. Gross, and A. Palomino, "Bayesian updating applied to the SHARP ATR system," in *Proceedings of the SPIE AeroSense 2000 Conference on Algorithms for Synthetic Aperture Radar Imagery VII*, **4053**, 2000.