Feature-aided Random Set Tracking on a Road Constrained Network

May 9, 2006

James Witkoskie, Walter Kuklinski, Stephen Theophanis, David Stein, and Michael Otero
MITRE Corporation
M/S E050
202 Burlington Road
Bedford, MA 01730-1420 U.S.A.
jwitkoskie@mitre.org

ABSTRACT

This paper describes the application of finite set statistics (FISST) to a real-time multiple target road constrained feature-aided tracking problem. A vehicle of interest traverses the road network while other confuser vehicles cross paths with this vehicle. Features extracted from sensors are used to disambiguate the vehicle of interest from the confuser vehicles. The FISST formalism naturally leads to understanding ambiguity in the identity of targets.

1.0 Introduction

Many tracking applications address the issue of monitoring one (or a few) vehicles of interest (VOI) in a background of confuser vehicles. One such program is the DARPA Affordable Moving Surface Target Engagement (AMSTE) program [1, 2]. Monitoring vehicles traversing a dense road network poses several complications since target state vectors can change rapidly due to target maneuvers at intersections [1]. There is often an association ambiguity between the measurement and the target positions because multiple roads and targets are in the sensor's field of view. Adding false alarms and missed detection further increase the complexity of road tracking. Since these confuser vehicles can cross paths with the VOI, ambiguity in the kinematic state of the VOI is unavoidable [2, 3]. Disambiguating the VOI from the confuser vehicle requires an ID process that verifies the position of the VOI using features that are unique to the VOI.

Traditional Kalman filter based approaches using multiple target tracking methods have difficultly in kinematic tracking of a VOI in a background of confuser vehicles because of limitations in the association process. The association process makes one or a few hard associations between targets and detections [4]. If incorrect, these hard associations lead to a greater certainty in target position than the data actually supports. The result is a false continuity in the track that implies known target IDs. This overly optimistic tracking result must be corrected by an algorithm introduced on top of the filter that accounts for confusion between tracks.

We have introduced and tested a random set tracker (RST) that naturally avoids making these associations in purely kinematic applications [5]. The random set approaches do not explicitly define tracks and avoid the association ambiguity by statistically weighing all possible hypotheses and associations [5, 6,7]. As a result, it does not have the implicit track continuity that leads to difficulties in defining track ambiguity. By Kronecker producting the kinematic space with a feature space, the RST incorporates ID into tracking problems without modifying the framework or introducing patches on top of an inadequate framework. Similar to its tracking capabilities, the RST framework naturally accounts for ID ambiguity by avoiding hard associations between feature vectors and specific tracks.

There are two basic feature-aided tracking applications that the current fusion algorithm can accommodate both applications [4]: 1) The simplest is the tracking vehicles of various classes, pick-up trucks versus sedans. In this application, the features of the various classes are pre-defined. Given a sensor measurement, the probability of a vehicle belonging to a specific class can be determined. 2) Another application requires that the system acquire features that uniquely identify the vehicle, such as high range resolution radar profiles [2]. After these features are acquired, the tracking application reduces to the classification problem with the assumption that only one vehicle belongs to each class. Determining when and how to acquire features is an inherent resource management issue. This paper outlines utilization of incoming sensor measurements to maintain tracks without explicitly addressing the issue of sensor management. Once the fusion framework is developed, it will be possible to predict future system behaviors, which will allow future development of resource management algorithms. The proposed RST implementation represents information in a framework that accommodates resource management algorithms performing both feature-aided tracking applications.

The specific application we discuss in this paper assumes the existence of two classes of targets, a VOI class and confuser class. As in the AMSTE program, we will assume the VOI class contains a unique target that is defined by its feature vector (FV), but unlike AMSTE, the FV of the VOI is known a priori. The FV space we consider is a continuous vector space, such as length or color, but the algorithm is easily modified to include a discrete FV. Measurements are used to refine our estimates of these vectors. The values of the FVs, along with kinematic discriminators such as speeds, are used to classify the vehicles into the VOI/confuser classes through traditional statistical classification processes, such as a support vector machine (SVM). The proposed fusion algorithm is not limited to this scenario, but this simplest scenario allows us to discuss several of the properties of this fusion algorithm.

2.0 Random Set Theory Tracking (RST)

Traditional Kalman filters assume state variables, $\mathbf{x}$, and measurements, $\mathbf{y}$, are fixed length random vectors [6]. For traditional tracking applications, $\mathbf{x}$ represents the targets' geokinetic variables, and $\mathbf{y}$ represents the measurements related to the geokinetic variables. The state vector motion model is linear with Gaussian white noise, $\mathbf{x}_{t+1} = \mathbf{\Omega}\mathbf{x}_t + d\boldsymbol{\lambda}_t$. Additionally, the measurements depend linearly on the state vectors with additive Gaussian white noise, $\mathbf{y}_t = \mathbf{W}\mathbf{x}_t + d\boldsymbol{\delta}_t$. The Kalman filter can be extended to non-Gaussian noise, non-linear measurements, and non-linear motion models through the Bayesian filter [6]. The Bayesian filter consists of two steps. 1) Starting with a probability density conditioned on previous measurements, $f(\mathbf{x}_{t-1} \mid \mathbf{y}_{1:t-1})$, a prediction step estimates the probability density of the vector at time step t,

$$f(\mathbf{x}_t \mid \mathbf{y}_{1:t-1}) = \int d\mathbf{x}_{t-1}\, f(\mathbf{x}_t \mid \mathbf{x}_{t-1})\, f(\mathbf{x}_{t-1} \mid \mathbf{y}_{1:t-1}). \tag{1}$$

This new probability density is combined with a new set of measurements, $\mathbf{y}_t$, during the update step to determine the new estimate of the probability density,

$$f(\mathbf{x}_t \mid \mathbf{y}_{1:t}) = N^{-1} f(\mathbf{y}_t \mid \mathbf{x}_t)\, f(\mathbf{x}_t \mid \mathbf{y}_{1:t-1}), \tag{2}$$

where N is a normalization factor. The number of state variables and measurements is fixed and the mapping between the state variables and measurements is explicit. The state variables, $\mathbf{x}$, can be expanded to include ID variables, FVs, such as length or color. These feature vectors correspond to characteristics of the vehicles being tracked. These FV variables allow us to discriminate targets that cannot be separated by kinematics alone. Equation (1) and (2) do not need to be modified to incorporate either discrete or continuous features vectors. Similar to the kinematic variables, FVs exist in a multidimensional space. Most collected features are time invariant and the feature's motion model is generally the identity operator, $f(\mathbf{x}_{t-1}) = f(\mathbf{x}_t)$, while measurements reduce our uncertainty in their values. Incorporating a time varying feature, such as temperature, is trivial.

Classification or identification, which generally results in a discrete set of states, is derived from the FVs through a partitioning of the FV space. Integration (marginalization) of the probability distribution over the interior of the classification boundaries defines the probability of the target belonging to a certain class (see figure 1). Defining the classification boundaries can be accomplished by several methodologies, including SVM. The hard decision boundaries can be replaced by the probability of a feature vector corresponding to a vehicle in a specific class. In this case, the classification is determined by $P_{class} = \int d\mathbf{x} f(\mathbf{x}) P_{class}(\mathbf{x}) \Big/ \sum_{class} \int d\mathbf{x} f(\mathbf{x}) P_{class}(\mathbf{x})$.

(3)

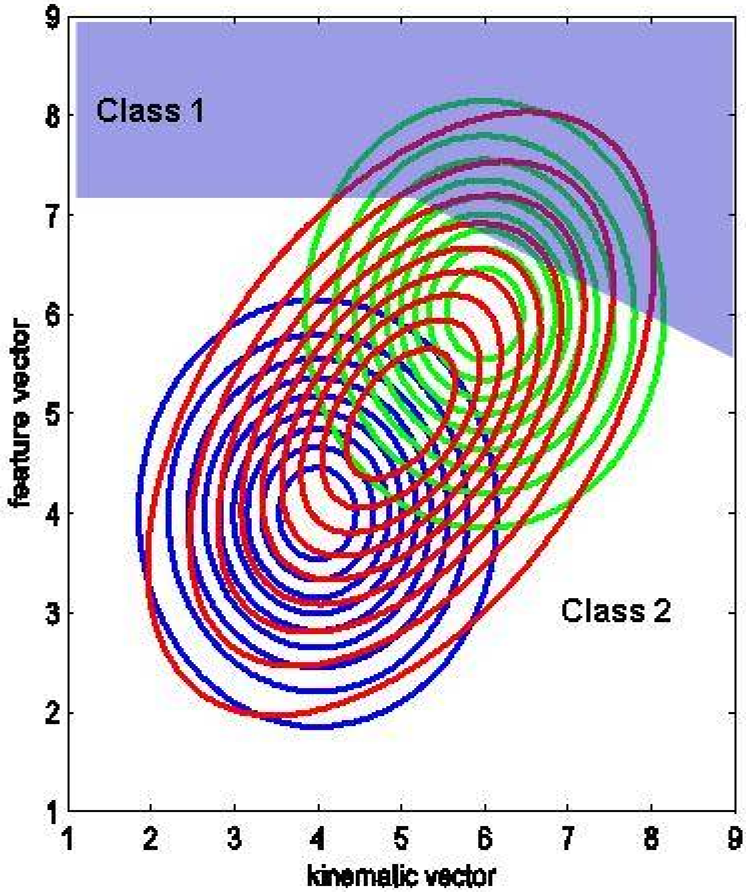

Figure 1: A depiction of object classification. The integral of the pdfs over the shaded portion of the vector space corresponds to the probability of being in class 1. The integral of the pdfs over the unshaded space corresponds to the probability of being in class 2. Note that the classification may depend on the kinematic variables. The green pdf straddles the classification boundary and exists in an ambiguous classification state. As discussed in section 4.0, the red pdf represents a hypothesis reduction that approximates the information contained in the the green and blue hypotheses. The ambiguity in which hypothesis is true becomes a larger variance in the pdf.

If $P_{class}(\mathbf{x})$ are mutually exclusive indicator functions, $P_{class}(\mathbf{x}) = 0,1$ and $P_{class1}(\mathbf{x})P_{class2}(\mathbf{x}) = \delta_{class1,class2}$, the hard decision boundaries result. The same arguments apply to a discrete feature vector space. In many ways the discrete space is a simplification of the continuous framework and will not be directly addressed in this paper.

Tracking multiple targets presents difficulties because the association of measurements with existing targets is often ambiguous [6,7]. Even if the state variables are linearly related to measurements with Gaussian errors, the ambiguity in associations produces non-Gaussian effects [3]. Missed detections, false alarms or clutter, and the birth and death of targets complicate the scenario. In addition to their numerical values, the number of targets and measurements are also random variables, and the vectors $\mathbf{x}$ and $\mathbf{y}$ must be replaced with random sets $\{x\}$ and $\{y\}$. The state variables, $\{x\}$, may take on the values $\{x\} = \{\phi\}$, $\{x^1, v^1, id^1\}$, $\{x^1, v^1, id^1, x^2, v^2, id^2\}, \ldots$ where x^i, v^i, id^i denote the positions, velocities, and FV components of target i, respectively. The set of measurements, $\{y\}$, are estimates of the geokinetic and feature variables recorded from various sensors and include clutter returns and missed detections.

Several methods address variable numbers of targets and detections with ambiguous associations, including joint integrated probabilistic data association (JIPDA) and jump Markov models (JMM) [10,11]. In this paper, we explore applications of finite set statistics (FISST) to road constrained multiple target tracking. FISST is a generalization of the Bayesian equations to sets, equations (1) and (2). The probability density, called the global density, is defined on the possible number and locations of targets. For road networks, the global density has the form

$$f_g = \begin{cases} f_0 \\ f_r(x,v,id) \stackrel{\Delta}{=} f_r(\{x^{(1)}\}) \\ f_{r_1 r_2}(x^1, v^1, id^1, x^2, v^2, id^2) \stackrel{\Delta}{=} f_{r_1 r_2}(\{x^{(2)}\}) \\ \vdots \end{cases}, \tag{4}$$

where f_0 is the probability of no targets, $f_{r_1 \ldots r_n}$ is the probability density for n targets on roads $r_1 \ldots r_n$ at $\{x^1, v^1, id^1, \ldots, x^n, v^n, id^n\} = \{x^{(n)}\}$, respectively. The density is normalized with respect to a set integral

$$\begin{aligned} &\sum_{n=0}^{\infty} \frac{1}{n!} \int f_{r_1 \ldots r_n}(x^1, v^1, \ldots, x^n, v^n) dx^1 dv^1 d(id^1) \ldots dx^n dv^n d(id^n) \\ &\quad = \sum_{n=0}^{\infty} \frac{1}{n!} \int f_{r_1 \ldots r_n}(\{x^{(n)}\}) d\{x^{(n)}\} = 1, \end{aligned} \tag{5}$$

where the summation over n refers to the number and identity of possible roads. Similar to the Bayesian filter one can define conditional probabilities, motion models, and measurement models to develop a set of recursive update equations [3, 7]. The prediction step includes propagation and birth and death processes so that the conditional expectation of the global density on previous measurements has the form

$$\begin{aligned} &f_{r_1 \ldots r_n}(\{x_t^{(n)}\} \mid \{y_{1:t-1}\}) = \\ &\sum_{m=0}^{\infty} \frac{1}{m!} \int d\{x_{t-1}\} f_{r_1 \ldots r_n | r_1' \ldots r_m'}(\{x_t^{(n)}\} \mid \{x_{t-1}^{(m)}\}) \; f_{r_1' \ldots r_m'}(\{x_{t-1}^{(m)}\} \mid \{y_{1:t-1}\}) \end{aligned} \tag{6}$$

where the sum is over all possible roads, $r_1' \ldots r_m'$. The predicted density is then updated with the measurements at time step t,

$$\begin{aligned} &f_{r_1 \ldots r_n}(\{x_t^{(n)}\} \mid \{y_{1:t}\}) = N^{-1} f_{r_1 \ldots r_n}(\{y_t^{(n)}\} \mid \{x_t^{(n)}\}) \\ &\quad \times f_{r_1 \ldots r_n}(\{x_{t-1}^{(n)}\} \mid \{y_{1:t-1}\}). \end{aligned} \tag{7}$$

Summation over the associations makes these expressions more complicated than equations (1) and (2). The exact forms of the expressions for the sensors used in this experiment will be outlined below. Similar to the Bayesian filter, the FISST tracker can be easily extended to applications with target ID by taking the Kronecker product of the FV space with the kinematic space. The FV extension of this previously proposed tracker handles the ambiguities in FV space in a manner that is analogous to the kinematic trackers handling of ambiguities.

3.0 Global Density Calculation

The prediction and update equations for the road constrained network are complicated by the association ambiguity of the sensors. In this section, we outline the expressions for these equations.

3.1. Prediction

The road networks we consider consist of road segments joined at intersections. Vehicles can enter and leave the road network via certain road segments that terminate outside the network. These terminated roads are referred to as exit ends. The prediction step includes propagation along a road, road switching, and birth and death processes. The centerline of each road, i, is defined by a parametric curve, $x_i(s)$ and $y_i(s)$, with $\sqrt{(dx_i/ds)^2 + (dy_i/ds)^2} = 1$. The independent motion of targets along the arc length of a road is linear with Gaussian white noise acceleration,

$$\begin{aligned} \dot{s}^i &= v^i \\ \dot{v}^i &= -\gamma_i v^i + d\Lambda_i \end{aligned} \quad . \tag{8}$$

Over a small time interval, the expected variance in the integral of the acceleration is $\left\langle \left(\int_0^{\Delta t} d\Lambda_i \right)^2 \right\rangle = \sigma_\Lambda^2 \Delta t$. The model's steady state root mean square velocity is $\sigma_{v_{ss}} = \sqrt{\sigma_\Lambda^2 / 2\gamma_i}$.

It is possible to include a FV dependence in the motion of the target. Target hypothesis branching at intersections is accomplished by allowing density to propagate beyond the end of the road. The overhanging density is truncated from the original road and equipartitioned among the connected roads (See Fig. (2)). The velocity profile of the equipartitioned distribution is the same as the original density, with velocities pointing towards the intersection on the original road pointing away from the intersection on the new road.

Death processes correspond to density that overhangs at exit ends in the road network. Death processes result in marginalizing over the overhanging profile and adding the remaining density profile to the corresponding hypothesis. For example, if Fig. (2) (b) corresponds to the projection of $f_{r_1 r_2}$ onto the r_1 coordinate at the exit end, the overhang would be truncated and $\sum_{id^1} \int_{x^1 > 10} dx^1 dv^1 f_{r_1 r_2}(x^1, v^1, id^1, x^2, v^2, id^2)$ is added to f_{r_2}. In the application we consider, the identity of targets does not vary over time and the FV component will not be affected by the target motion. These assumptions can be easily modified.

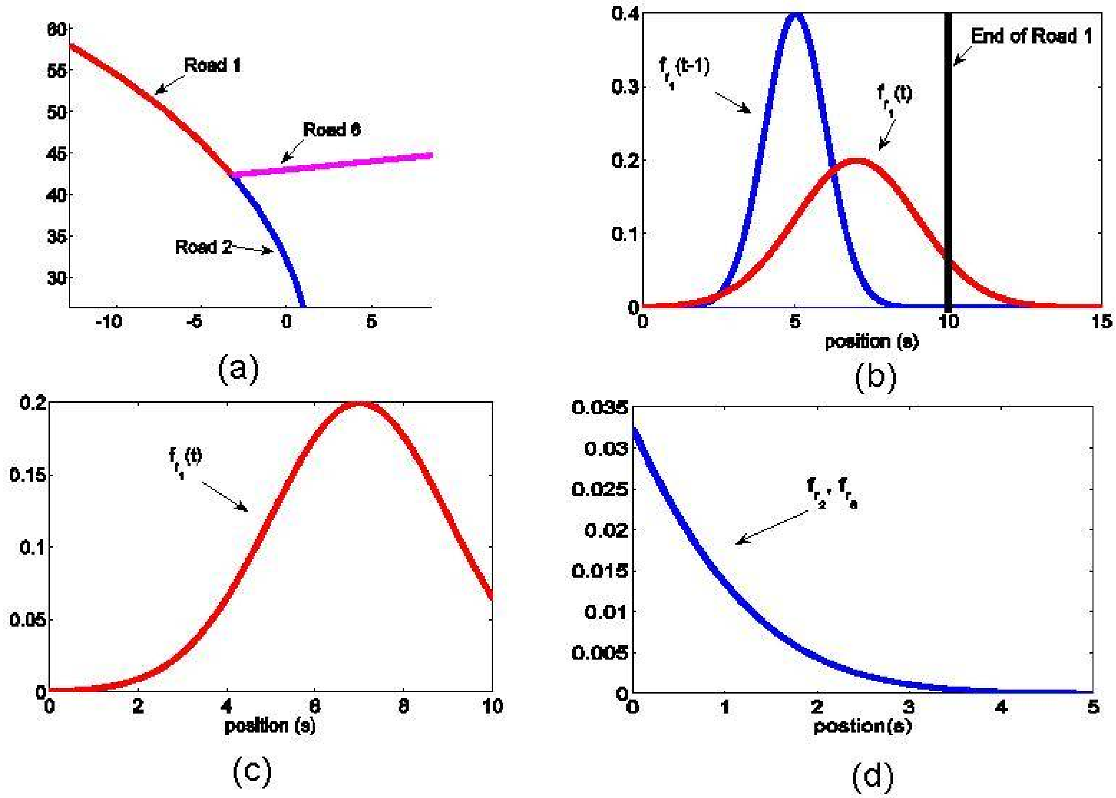

Figure 2: Road switching in the road network. (a) The intersection of r_1, r_2 and r_6. (b) Starting from a density on r_1, the propagation step results in the density overhanging r_1. (c) The new profile for r_1 is created by truncating the overhanging portion of the density. (d) The overhanging density is equipartitioned between r_2 and r_6.

Several birth processes are possible, such as coupling birth processes to measurements not associated with any previously detected targets [7]. Resource management of sensors requires knowledge of the rate of decay of information about the global density, which is greatly affected by birth processes. To ensure that the birth processes reflect the decay in information, the birth processes correspond to adding targets to an exit end by a Poisson process with parameter, $\lambda_{r_i} << 1$. The added target has a Gaussian position profile with a mean located near the end of the road and a road dependent standard deviation. The velocity profile is a truncation of the steady state velocity profile that ensures the velocity points onto the road so that targets do not immediately depart upon entry. Birth targets will generally exist in an ambiguous classification state with a widely distributed FV that straddles both the VOI and confuser classifications (see figure 1).

3.2. Update

A sensor is an instrument that returns detections based on data collected from the environment. The detections can contain both kinematic information and FV information. The kinematic information will generally be in the form of range, azimuth, and elevation position and velocity information that is relative to the position of the sensor. Two examples are a range radar that reports only a range from the radar to the target and an acoustic array that reports an angle from the sensor to the target. Every sensor has a probability of detecting a target and errors in reported measurements. The probability of detect and the errors in the measurement depend on the targets locations relative to the sensor and may even contain multi-body effects, such as occlusion of one target by another.

As mentioned above, summation over the associations complicates updating the probability density. For example, if a radar sensor reports the range measurement, $\{\rho_1 ... \rho_m\}$, the likelihood function becomes

$$f_{r_1 ... r_n}(\{\rho_1 ... \rho_m\} \mid \{x^{(n)}\}) = \sum_{\substack{\rho_1 = (clutter, 1...n) \\ \vdots \\ \rho_m = (clutter, 1...n)}} \left[\prod_{\substack{b \neq \rho_a \\ \rho_a = clutter}} (1 - P_c(b)) \prod_{\rho_a = clutter} P_c(\rho_a) \prod_{t_k \neq \rho_a} (1 - P_t(t_k)) \right. \tag{9}$$

$$\left. \times \prod_{t_k = \rho_a} P_t(t_k) g_r(\rho_a - h_r(x^k(s^k), y^k(s^k))) \right].$$

In this expression, ρ_a is the measurement (in the product and summation) and the association (in the indices). The probability of a clutter return from bin b, $P_c(b)$, or from the target, P_t, is spatially dependent. The measurement model, g_r, is often Normally distributed error around the point in the measurement space, $h_r(x(s), y(s)) = \sqrt{x(s)^2 + y(s)^2}$, which is a non-linear map between the state space and measurement space.

Association ambiguities lead to non-Gaussian effects and make the Kalman filtering equations invalid even if h_r was linear, g_r is Normally distributed, and the detection probabilities are position independent constants. The non-linear nature of the measurements and detection probabilities increases these complications. Various linear approximations, pruning, and measurement weighting strategies attempt to avoid this difficulty, but they may lead to poor tracking performance [3, 10].

The association between a detection from a sensor and a target depends on both the sensor and the targets. As an example, shot locating acoustic arrays have a markedly different association algorithm than radars. The array often reports a single measurement, the angle to the loudest target in the sensors field of view (the direction to the shooter), although more complex behavior can sometimes result. There are fewer associations to sum over, but the detection probabilities will depend on all coordinates of all possible shooters. Neglecting elevation, the likelihood of detecting an azimuth of α is

$$\begin{aligned} & f_{r_1 ... r_n}(\{\phi\} \mid \{x^{(n)}\}) = P(T > clutter, s_j) \\ & f_{r_1 ... r_n}(\{\alpha\} \mid \{x^{(n)}\}) = P(clutter > T, s_j) + \\ & \quad \sum_{s_i = t_1 ... t_n} P(s_i > clutter, T, s_{j \neq i}) g_\alpha(\alpha - h_\alpha(x^i)) \end{aligned} \tag{10}$$

where s_j refers to the signal strength from the targets, clutter refers to background sound sources, and T is a threshold for reporting detections. Similar to the radar model, g_α will typically be Normally distributed and the relation between road positions and angle, $h_\alpha(x(s), y(s)) = \tan^{-1}(y(s)/x(s))$, is highly non-linear. These types of instruments require many-body effects that cannot be captured by a PHD filter [12].

Variations transverse to the centerline of the road are also utilized in combining measurements with the priors. These variations are assumed to be proportional to the width of the road. The transverse variations allow the system to incorporate detections that are slightly off the road. If the roads have significant changes in elevation, this component could be easily incorporated into these models by defining $z(s)$.

4.0 Numerical Implementation

To make the random set implementation numerically feasible, we use a Gaussian Mixture approximation. The Gaussian mixture model for this road network scenario also allows a scalable Gaussian sum particle filter (GSPF) representation [13]. The GSPF is similar to particle filter sampling methods, but the delta function kernel associated with each particle is replaced by a variable dimensional Gaussian. The covariance and mean of each particle is propagated instead of simply the position of each particle. Each term in the global density is represented by a finite number of Gaussians,

$$f_{r_1 \dots r_n}(\{x^{(n)}\}) = \sum_i a_i G(\mathbf{x}^{(n)} - \boldsymbol{\mu}_i^{(n)}, \mathbf{C}_i^{(n)}) \tag{11}$$

where G is a multidimensional Gaussian distribution, $G(\mathbf{x}, \mathbf{C}) = \sqrt{Det(2\pi\mathbf{C})}^{-1} \exp(\mathbf{x}^T \mathbf{C}^{-1} \mathbf{x})$ which includes the FV. The Gaussian mixture representation requires several approximations: 1) After combining detections with the Gaussian, the probability of detect is based on the mean value of each Gaussian component. 2) The non-linear maps between the coordinates and the measurements are Taylor expanded around the mean values, $h \approx h(x(s_0)) + (\partial h(x(s_0))/\partial s)(s - s_0)$ in directions parallel and transverse to the road network. 3) The motion model is still linear. 4) Instead of truncating overhanging density, road switching and death processes correspond to the mean of the mixture component overhanging the end of the road. The entire Gaussian component either switches roads or is marginalized. 5) Birth processes correspond to adding mixture components to the ends of roads with a mean velocity pointing onto the roads and a fixed initial width.

All of these approximations maintain the Gaussian mixture representation of the global density. The validity of these approximations depends on the Gaussian mixture components' variances being much smaller than variations in the detection probabilities, the variations in h, and the distance from the ends of the roads. We ensure the validity of these approximations by replacing components with large variances with several Gaussian components of smaller variances through a Kullback-Leibler measure. Similarly, if a target is on the boundary of a classification, it is possible to split up a Gaussian into multiple Gaussian mixtures that do not straddle the boundary, but we have not implemented this approach [13].

Although these approximations ensure that the global density maintains a Gaussian mixture functional form, the associations of the measurements result in a geometric explosion in the number of mixture components. To avoid the geometric growth, the mixture components are recombined based on the Kullback-Leibler metric. The KL metric compares $f_{r_1 \ldots r_n}$ with a new Gaussian mixture, $\tilde{f}_{r_1 \ldots r_n}$, that contains two components with the same mean, variance and class probability [14],

$$\begin{aligned}\tilde{f}_{r_1 \ldots r_n}(\{x^{(n)}\}) &= \sum_k b_k G(\mathbf{x}^{(n)} - \mathbf{m}_k^{(n)}, \mathbf{M}_k^{(n)}) = \sum_{k=i,j} a_k G(\mathbf{x}^{(n)} - \tilde{\boldsymbol{\mu}}_k^{(n)}, \tilde{\mathbf{C}}_k^{(n)}) \\ &+ \sum_{k \neq i,j} a_k G(\mathbf{x}^{(n)} - \boldsymbol{\mu}_k^{(n)}, \mathbf{C}_k^{(n)}).\end{aligned} \tag{12}$$

The KL metric gives

$$\begin{aligned}D_{f|\tilde{f}} &= \sum_k \int d\{x^{(n)}\} a_k G(\mathbf{x}^{(n)} - \boldsymbol{\mu}_k^{(n)}, \mathbf{C}_k^{(n)}) \ln\left(\frac{a_k G(\mathbf{x}^{(n)} - \boldsymbol{\mu}_k^{(n)}, \mathbf{C}_k^{(n)})}{b_k G(\mathbf{x}^{(n)} - \mathbf{m}_k^{(n)}, \mathbf{M}_k^{(n)})} \right) = \\ &\sum_{k=i,j} \int d\{x^{(n)}\} a_k \left[G(\mathbf{x}^{(n)} - \boldsymbol{\mu}_k^{(n)}, \mathbf{C}_k^{(n)}) \ln\left(\frac{G(\mathbf{x}^{(n)} - \boldsymbol{\mu}_k^{(n)}, \mathbf{C}_k^{(n)})}{G(\mathbf{x}^{(n)} - \tilde{\boldsymbol{\mu}}_k^{(n)}, \tilde{\mathbf{C}}_k^{(n)})} \right) \right].\end{aligned} \tag{13}$$

where $\tilde{\boldsymbol{\mu}}^{(n)}$ and $\tilde{\mathbf{C}}^{(n)}$ are the optimal mean and covariance, respectively. The combined mixture components that result in the smallest KL metric are combined until either the number of mixtures is smaller than a set maximum number or the KL metric for combining two mixtures is greater than a tolerance [14]. The procedure results in a reduction in the number of mixture components, as demonstrated in figure 1.

If the FV does not affect the kinematic model we may want to ignore cross terms in the information and evaluate the information about the feature and kinematic variables separately. In this case, the KL metric will factorize into a FV component and a kinematic component.

$$\begin{aligned}D_{f|\tilde{f}} &= \sum_{k=i,j} \int d\{x^{(n)}\} a_k G(\mathbf{x}^{(n)} - \boldsymbol{\mu}_k^{(n)}, \mathbf{C}_k^{(n)}) \\ &\times \left[\ln\left(\frac{G(\mathbf{x}_{kin}^{(n)} - (\boldsymbol{\mu}_{kin}^{(n)})_k, (\mathbf{C}_{kin}^{(n)})_k)}{G(\mathbf{x}_{kin}^{(n)} - \tilde{\boldsymbol{\mu}}_{kin}^{(n)}, \tilde{\mathbf{C}}_{kin}^{(n)})} \right) + \alpha \ln\left(\frac{G(\mathbf{x}_{id.}^{(n)} - (\boldsymbol{\mu}_{id}^{(n)})_k, (\mathbf{C}_{id}^{(n)})_k)}{G(\mathbf{x}_{id}^{(n)} - \tilde{\boldsymbol{\mu}}_{id.}^{(n)}, \tilde{\mathbf{C}}_{id.}^{(n)})} \right) \right].\end{aligned} \tag{14}$$

The kin and id subscripts denote the kinematic and FV components of the Gaussian mixture, respectively. The α prefactor allows us to weigh the importance of the feature and kinematic information. It is also possible to give different geokinetic variables, such as those corresponding to VOI, different weights as well.

If the final distribution was reduced to a single Gaussian at each time step, the resulting filter would be similar to the JIPDA filter, which sums over associations to determine the best fitting Gaussian distribution [10]. Maintaining a finite number of mixture components gives greater flexibility and improves performance. Each mixture component is similar to a single hypothesis in a multiple hypothesis tracker (MHT) [15]. Unlike the MHT that performs branching and pruning procedures, the FISST tracker combines branches of the track estimates [14]. This procedure reduces the amount of information lost at each update step since the combined hypothesis contains information about the original hypotheses, but makes the definition of a track and track lifetime ambiguous since tracks with different lifetimes may be combined (see figure 3).

5.0 Discussion

The random set tracker proposed in this paper is closely related to multi-threaded Kalman filtering methodologies such as MHT, but there are several distinctions that naturally lead to incorporation of target identity. The linearization in measurement models is identical to the extended Kalman filter. In fact, if there were no association ambiguities, (i.e. we knew which target creates each detection), the resulting algorithm would be an extended Kalman filter. Similarly, if we always reduced the distribution to a single Gaussian, a tracker similar to the JIPDA filter would result [10]. However, the random set tracker sums over associations, which avoids hard decisions made in many traditional Kalman Filter based approaches. A MHT attempts to track several possible hard associations, but introduces a pruning procedure to reduce the number of possible hypotheses. Unlike MHT, the random set tracker avoids the exponential branching of hypotheses by reducing information (not the explicit number of hypotheses). The number of hypotheses is allowed to vary as needed to maintain the appropriate level of information about the targets. Figure 3 gives an illustration of the differences between MHT and RST. Figure 3 (a) shows a path where a confuser vehicle will become confused with the VOI track at two points in time (1 and 3). The multiple hypothesis tracker will have two branching processes, resulting in four hypotheses. If the number of hypotheses maintained by the tracker is less than four, some of these hypotheses will be eliminated despite that the probability of each hypothesis may be close to equally likely. The RST tracker would prevent this branching and pruning by determining that the hypotheses at time point 3 contain similar information and can be approximated by a single Gaussian with little loss of information. This reduces the number of hypotheses handled by the tracker. When the second branch occurs, only two hypotheses result.

Another example of the RST is a scenario where there are two hypotheses of a single target with distinct FVs. This scenario can result if the kinematic detections suggest the existence of a single target, but FVs taken from different sensors are not complimentary. In this case, it is most probable that only a single FV is correct. Since we do not know the correct FV, our knowledge of the FV is diminished, which is captured by combining the Gaussians of the two FVs into one Gaussian with a larger variance (see figure 1). The MHT would either prune one hypothesis and lose information or keep both hypotheses, which increases the cost of maintaining this information. The RST's ability to incorporate this FV ambiguity while reducing the computation of the hypotheses is a advantageous property of RST.

Similar to MHT, the RST naturally deals with ambiguities in ID by leading to different hypotheses where the confuser and VOI position may be switched. However, these multiple hypotheses are not the result of a few possible hard associations. This scenario is quite different from traditional multiple thread Kalman filter algorithms that attach a track number and give a false continuity to the tracks. It is important to emphasize the difference between classification (which also attaches a label to a track) versus the track label in multiple threaded Kalman filtering. The classification label is not an intrinsic part of the fusion algorithm. Instead, it is a refined data product output by the fusion algorithm. This classification product would be important to a resource management algorithm or an end user, but it is not directly used by the fusion algorithm.

The RST fusion algorithm also lacks the traditional concept of a track, which is a contiguous set of position estimates over time believed to be the result of a single vehicle. Internally, the RST fusion algorithm simply reports the probability of a set of target with specific FVs and geokinetic variables at a specific time point. A continuous track can be inferred from a MAP estimate performed on the global density, but this track is a refined data product that cannot be used internally by the fusion algorithm.

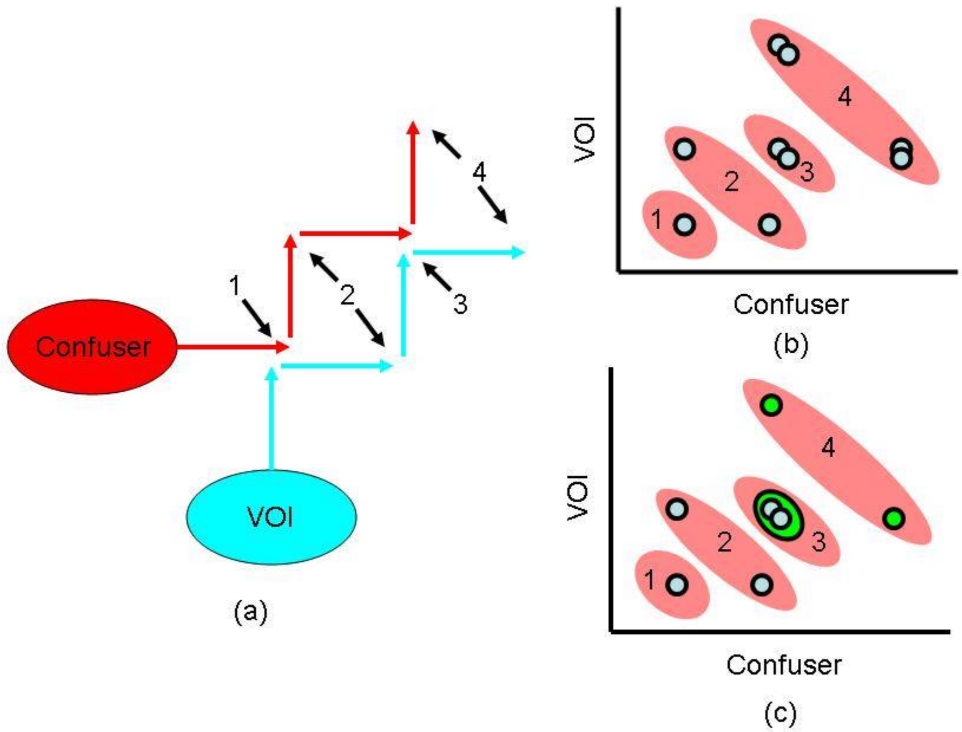

Figure 3: (a) A path where the VOI and confuser become entangled at points 1 and 3. If no ID measurements can be performed at point 2, a MHT will give 4 possible hypotheses at point 4 as shown in (b). The random set tracker recombines the two hypotheses at point 3 so that only two hypotheses exist at point 4.

The RST fusion algorithm is a promising candidate for feature-aided target tracking. It naturally addresses branching processes and identity ambiguity within a computable framework. We have currently implemented a RST on a small road network with 6 roads and 6 sensors including radars, acoustics and cameras. The tracker can work at real time tracking up to 3 targets at a 2 Hz rate on a standard laptop computer. Analysis of these results will be presented in future work. Extending this RST to larger applications requires development of algorithms such as windowing sensor detections and clustering of targets. Because the RST framework closely resembles a MHT, these algorithms can be easily incorporated into RST and pose a future research direction.

References

[1] T. Kirubarajan, Y. Bar-Shalom, K.R. Pattipati, I. Kardar, B. Abrams, and E. Eadan. "Tracking Ground Targets with Road Constraints Using an IMM Estimator." *IEEE Aerospace Conference Proceedings.* 5:5—12, Mar. 1998.

[2]D.R. Kirk, T. Grayson, D. Garren, and Chee-Yee Chong. "AMSTE Precision Fire Control Tracking Overview." *IEEE Aerospace Conference Proceedings.* 3:465—472, Mar. 2000.

[3] C.A. Scott and C.R. Drane. "An Optimal Map-Aided Position Estimator For Tracking Motor Vehicles." *IEEE Proceedings of the 6th International Conference on Vehicle Navigation and Information Systems.* 360--367, Jul. 1995.

[4] D.H. Nguyen, J. H. Kay, B. J. Orchard, and R. H. Whiting. "Feature-aided tracking of moving ground vehicles." *Proceedings of SPIE.* 4727:234—245, Aug. 2002

[5] J.B. Witkoskie, W.S. Kuklinski, S. Theophanis, M. Otero. "Random Set Tracker Experiment on a Road Constrained Network with Resource Management." *Proceedings of the 9th International Conference on Information Fusion.* In publication, 2006.

[6] I.R. Goodman, R.P.S. Mahler, and H.T. Nguyen. *Mathematics of Data Fusion*. Kluwer Academic Publishers. Dordrecht, Netherlands. 1997.

[7] M.R Morelande and S. Challa. "A Multitarget Tracking Algorithms Based on Random Sets." IEEE Proceedings of the 6th International Conference of Information Fusion. 2:807—814, 2003.

[8] B.N. Vo and S. Singh, and W.K. Ma. "Tracking Multiple Speakers Using Random Sets." *IEEE Conference Proceedings on Acoustics, Speech, and Signal Processing*. 2: ii357-ii360, 2004.

[9] R. Evans, V. Krishnamurthy, G. Nair, and L. Sciacca. "Networked Sensor Management and Data Rate Control for Tracking Maneuvering Targets." *IEEE Transactions on Signal Processing*. 53(6):1979-1991, 2005.

[10] D. Musicki and R. Evans. "Joint Integrated Probabilistic Data Association." *IEEE transactions on Aerospace and Electronic Systems*. 40(3): 1093—1099, 2004.

[11] A. Doucet, B.N. Vo, C. Andrieu, and M. Davy. "Particle Filtering for Multi-Target Tracking and Sensor Management." *Proceedings of the Fifth International Conference on Information Fusion*. 1:474—481, Jul. 2002.

[12] R.P.S. Mahler. "Multitarget Bayes Filtering via First-Order Multitarget Moments." *IEEE Transactions on Aerospace and Electronic Systems*. 39(4):1152—1178, Oct. 2003.

[13] J.H. Kotecha and P.M Djuric. "Gaussian Sum Particle Filtering." IEEE Transactions on Signal Processing. 51(10):2602—2612, Oct. 2003.

[14] W. Adam, R. Fruhwirth, A. Strandlie, and T. Todorov. "Reconstruction of elections with the Gaussian-sum filter in CMS tracker at LHC." *Journal of Physics G-Nuclear and Particle Physics*. 31:N9-N20, Mar. 2005.

[15] S.S. Blackman. "Multiple Hypothesis Tracking For Multiple Target Tracking." IEEE Aerospace and Electronic Systems Magazine. 19:5—18, 2004.

[16] P. Whittle. "Restless bandits: Activity allocation in a changing world." In *A Celebration of Applied Probability*. J. Gani (Ed.). *Journal of Applied Probability*. 25a: 287—298, 1988.

[17] D. Bertsimas and J. Nino-Mora. "Restless bandits, linear programming relaxations and a primal-dual index heuristic." *Operations Research*. 48(1):80—90, Jan. 2000.

[18] M. Coats. "Distributed Particle Filters for Sensor Networks." In *Proceedings of the 3rd International Symposium on Information Processing in Sensor Networks*. ACM Press. NY, NY. 99-107, Apr. 2004.