2008 American Control Conference
Westin Seattle Hotel, Seattle, Washington, USA
June 11-13, 2008

Data Association with Ambiguous Measurements

Matthew Travers, Todd Murphey, and Lucy Pao

***Abstract*— We address the problem of tracking a single object in the neighborhood of several other closely spaced, similar objects where the sensor used to do the tracking may randomly measure the wrong object. Unlike many tracking scenarios, there is no other environmental clutter producing additional erroneous measurements. The objects move together, and the sensor provides one measurement at every time step, either due to the object of interest or due to one of the other similar nearby objects. This situation of having a "mixed" set of measurements of unknown origin occurs in real world systems. While we consider the mixed-measurement problem in an example scenario, the algorithms developed can be applied to any number of associated systems with little alteration.**

I. Introduction

Tracking a single object in an unknown environment is a relatively well studied process. Tracking a group of objects in formation has also been studied by many. However, tracking a single object in a group has not.

An example problem of tracking a single object in a group is tracking a single automobile in a convoy, the example considered later in this paper. In the example, the cars all move along a known road together in formation. This makes the problem difficult because there are minimal differences between individual trajectories over the entire sampling period. Another aspect of the problem we will consider that adds difficulty is how the sensor makes measurements. The sensor only returns one measurement per time step, and it is not known whether the measurement comes from the particular car of interest or simply one of the other vehicles in the convoy. As will be shown, certain data association algorithms can be used to effectively sort out the measurements and provide reliable tracking.

The data association algorithms presented are not only applicable to systems with spatial distributions as considered here, but they are also applicable to systems with temporal distributions. An example of the latter occurs when a distributed network of Geiger counter sensors is used to track a moving nuclear source. This situation is similar to the convoy of cars, except that instead of a second nearby car producing erroneous measurements, background radiation acts to interfere with the measurements coming from the "object of interest."

The data association algorithms covered in this paper are the extended Kalman filer (EKF), gating, probabilistic data association (PDA), and a Kolmogorov-Smirnov (KS) test. The extended Kalman filter is a standard estimation algorithm for nonlinear systems. Gating is a method of eliminating measurements that are unlikely to have originated from the object of interest. PDA computes the probability of each measurement originating from the object of interest. The Kolmogorov-Smirnov test is a new method we are adapting for determining the origin of a particular measurement.

This work has been supported in part by the US Air Force Office of Scientific Research under Award No. FA9550-07-1-0311

Department of Electrical and Computer Engineering, University of Colorado at Boulder, Boulder, Colorado 80309

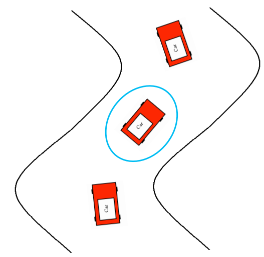

Fig. 1. Convoy of cars example.

We will present a new metric to be used in a Kolmogorov-Smirnov like test that provides better performance than the standard KS test metric in the tracking problem of interest.

This paper is organized as follows. An overview of the example convoy of cars problem is further developed in Section II, and Section III discusses data association algorithms to address the problem. The methods are applied to a convoy of cars example in Section IV, and simulation results are presented in Section V.

II. Example Problem Overview

Consider a convoy of cars moving along a known road. The convoy can consist of any number of vehicles. One of the cars is deemed to be of interest, i.e., there is a specific car that we would like to track. This situation is illustrated in Figure 1, where the car of interest is circled.

In this initial work, we assume that there is no maneuvering as the cars move along the road. The nonlinear model for each of the cars is

$$x_i(k) = f[k-1, x_i(k-1)] + v_i(k-1), i = 1, 2, ..., N \quad (1)$$

where $x_i(k)$ is the state at time k. N is the number of cars in the convoy. The additive process noise, $v_i(k-1)$, is zero-mean and white with covariance $Q_i(k-1)$. The measurement due to each vehicle is

$$z_i(k) = h[k, x_i(k-1)] + w_i(k-1), i = 1, 2, ..., N$$

where additive noise, $w_i(k-1)$, is also assumed to be zero-mean and white with covariance $R(k-1)$. It is assumed that the process noise, measurement noise, and initial state $\hat{x}(0|0)$ with the associated covariance $P(0|0)$ are all uncorrelated. The state $x \in \Re^n$ and the measurement $z \in \Re^{n_z}$, where n

is the dimension of the state space and n_z is the dimension of the measurement space.

The cars are assumed to be constrained to stay in a line on the road, and the individual separation between cars remains on average constant. The road is assumed to be perfectly flat and a model is defined for it.

The sensor is fixed to a tower with an unobstructed view of the entire road where the vehicles are traveling and is tasked to always be pointing in the direction of the car of interest. The sensor only makes measurements of Cartesian position, in a frame fixed to the base of the sensor tower. A single measurement is taken at each time step and the origin of that measurement is not known. For example, consider two cars traveling in a convoy, where we are trying to track car 1. Five measurements are taken, one of which is randomly from car 2. One possibility for how the list of measurements recorded by the sensor appears is

$$\begin{aligned}&\{z(k), z(k+1), z(k+2), z(k+3), z(k+4)\} = \\ &\quad \{z_1(k), z_1(k+1), z_1(k+2), z_2(k+3), z_1(k+4)\}.\end{aligned}$$

Only the set of measurements (with unknown origin) is provided. The actual identity of individual measurements on the right hand side is not known. The data association algorithm must sort out which measurement in the set originated from the car of interest.

III. ALGORITHMS

A. The Extended Kalman Filter

The EKF is a well known data association method for filtering nonlinear systems [2]. The EKF is an estimation algorithm that maps the current estimate of the state, $\hat{x}(k-1|k-1)$, forward one time step. The current estimate is conditioned on the current state given the current measurement [2], i.e., $\hat{x}(k-1|k-1) \triangleq E[x(k-1)|Z^{k-1}]$ where $Z^{k-1} \triangleq \{z(j), j = 1, \ldots, k-1\}$ is the cumulative set of measurements up to time $k-1$. The associated state error covariance matrix is $P(k-1|k-1) \triangleq E\{[x(k-1)-\hat{x}(k-1|k-1)][x(k-1)-\hat{x}(k-1|k-1)]'|Z^{k-1}\}$.

The prediction $\hat{x}(k|k-1)$ is computed by expanding the nonlinear function (1). The expansion is accomplished by evaluating the Taylor series of (1) about the current estimate $\hat{x}(k-1|k-1)$,

$$\begin{aligned}x(k) = f[k-1, &\hat{x}(k-1|k-1)] \\ &+ f_x(k-1)[x(k-1)-\hat{x}(k-1|k-1)] \\ &+ 1/2\sum_{i=1}^{n_x} e_i[x(k-1)-\hat{x}(k-1|k-1)]' \\ &\cdot f_{xx}^i(k-1)[x(k-1)-\hat{x}(k-1|k-1)] \\ &\qquad + \text{higher}-\text{order terms} + v(k-1) \quad (2)\end{aligned}$$

where e_i is the i^{th} Cartesian basis vector and $f_x(k-1)$ is the Jacobian of the vector f evaluated at time k, and $f_{xx}^i(k-1)$ is the Hessian of the i^{th} component of f.

The prediction of the state at time k is then obtained by taking the expectation of the Taylor series expansion (2) conditioned on Z^{k-1} and neglecting higher-order terms:

$$\begin{aligned}\hat{x}(k|k-1) &= f[k-1, \hat{x}(k-1|k-1)] \\ &\quad + 1/2\sum_{i=1}^{n_x} e_i \mathrm{tr}[\mathrm{f}_{\mathrm{xx}}^{\mathrm{i}}(\mathrm{k}-1)\mathrm{P}(\mathrm{k}-1|\mathrm{k}-1)].\end{aligned}$$

The state prediction error is

$$\begin{aligned}&\tilde{x}(k|k-1) = f_x(k-1)\tilde{x}(k-1|k-1) \\ &+ 1/2\sum_{i=1}^{n_x} e_i[\tilde{x}'(k-1|k-1)f_{xx}^i(k-1)\tilde{x}(k-1|k-1) \\ &\qquad - \mathrm{tr}[\mathrm{f}_{\mathrm{xx}}^{\mathrm{i}}(\mathrm{k}-1)\mathrm{P}(\mathrm{k}-1|\mathrm{k}-1)]] + \mathrm{v}(\mathrm{k}-1)\end{aligned}$$

and its covariance is

$$\begin{aligned}P(k|k-1) &\triangleq E[\tilde{x}(k|k-1)\tilde{x}'(k|k-1)|Z^{k-1}] \\ &= f_x(k-1)P(k-1|k-1)f_x' \\ &+ 1/2\sum_{i=1}^{n_x}\sum_{j=1}^{n_x} e_i e_j' tr[f_{xx}^i(k-1) \\ &\cdot P(k-1|k-1)f_{xx}^j(k-1)P(k-1|k-1)] \\ &\qquad + Q(k-1).\end{aligned}$$

The measurement prediction is

$$\begin{aligned}\hat{z}(k|k-1) &= h[k, \hat{x}(k|k-1)] \\ &\quad + 1/2\sum_{i=1}^{n_z} e_i \mathrm{tr}[\mathrm{h}_{\mathrm{xx}}^{\mathrm{i}}(\mathrm{k})\mathrm{P}(\mathrm{k}|\mathrm{k}-1)].\end{aligned}$$

where $h_{xx}^i(k)$ is the Hessian of the i^{th} component of h. The innovation is

$$\nu(k) = z(k) - \hat{z}(k|k-1), \quad (3)$$

and the innovation's covariance is

$$\begin{aligned}S(k) &\triangleq h_x(k)P(k|k-1)h_x'(k) \\ &+ 1/2\sum_{i=1}^{n_z}\sum_{j=1}^{n_z} e_i e_j'[h_{xx}^i(k)P(k|k-1) \\ &\cdot h_{xx}^j(k)P(k|k-1)] + R(k)\end{aligned}$$

where $h_x(k)$ is the Jacobian of h evaluated at time k. The EKF gain is $W(k) \triangleq P(k|k-1)h_x'(k)S^{-1}(k)$. The filtered estimate is then $\hat{x}(k|k) = \hat{x}(k|k-1) + W(k)\nu(k)$ and the filtered estimate error covariance is $P(k|k) = [I - W(k)h_x(k)]P(k|k-1)$.

B. Gating

Gating determines whether it is likely that a particular measurement came from the object being tracked. Assuming the measurement at each time step is normally distributed around the expected measurement, it is possible to define a region in the measurement space where there is a high probability of finding the measurement [2]

$$\tilde{V}_k(\gamma) \triangleq \{z : \nu'(k)S^{-1}(k)\nu(k) \le \gamma\}, \quad (4)$$

where ν is the innovation as in (3). The probability of gating the true measurement is

$$P_G = \int_0^{\gamma} \frac{\gamma^{n_z/2-1}\exp\{-\gamma/2\}}{2^{n_z/2}\Gamma(n_z/2)} d\gamma, \gamma \geq 0$$

where $\Gamma(\cdot)$ is the gamma function and n_z is the dimension of the measurement space. Assuming the measurement error is Gaussian, this norm is chi-square distributed. If the measurement error is not Gaussian, such as when accounting for the underlying system geometry in the measurement, gating can still be carried out with some modification [10].

As defined in Sections I and II, at each time step there is exactly one measurement of unknown origin received. If the measurement falls within the gate, it is used in a filtering algorithm. If the measurement falls outside the gate, the filtering algorithm uses the system model alone to predict forward one time step.

C. Probabilistic Data Association (PDA)

PDA has been used when multiple measurements are received at each time step in tracking an object in a region with uniformly distributed clutter [2]. The probability β_i of each measurement z_i being the true measurement from the object of interest is computed. The PDA is used in conjunction with a filtering algorithm, such as the EKF, where updated estimates $\hat{x}_i(k|k)$ due to each measurement z_i are computed and the overall updated state estimate is

$$\hat{x}(k|k) = \sum_{i=0}^{m_k} \hat{x}_i(k|k)\beta_i(k), \tag{5}$$

where m_k is the number of measurements at time k, and $\sum_{i=0}^{m_k} \beta_i(k) = 1$, i.e., the probabilities are mutually exclusive as well as exhaustive. When tracking in cluttered environments where multiple measurements are received at each time step, gating is often used to limit the computational complexity. When gating is used, m_k in (5) is the number of gated measurements.

When tracking with some types of sensors (e.g., radar or sonar), the probability P_D of detecting the object of interest is often less than 1. The $i = 0$ case in (5) represents the hypothesis that no true measurement from the object of interest is received at time k. This implies that for the type of sensor described in Section II, if the one measurement per time step is determined to not be from the object of interest, the update (5) will be entirely based on the $i = 0$ case.

The error covariance associated with (5) is

$$P(k|k) = \beta_0(k)P(k|k-1) + [1 - \beta_0(k)]P^c(k|k) + \tilde{P}(k)$$

where $P(k|k-1)$ is the same as is in the EKF, $\tilde{P}(k) = W(k)[\sum_{i=1}^{m_k} \beta_i(k)\nu_i(k)\nu_i'(k) - \nu(k)\nu'(k)]W'(k)$ where $\nu_i(k) = z_i(k) - \hat{z}(k|k-1)$, $\nu(k) = \sum_{i=1}^{m_k} \beta_i(k)\nu_i(k)$, and $P^c(k|k) = [I - W(k)h_x(k-1)]P(k|k-1)$.

Clutter measurements are assumed to be uniformly distributed with density $\lambda = \frac{m_k}{V_k}$, where V_k is the volume of the validation region. $V_k = c_{n_z}\gamma^{n_z/2}|S(k)|^{1/2}$, and c_{n_z} is the volume of the n_z-dimensional unit hypersphere. The $\beta_i(k)$ probabilities are computed as [2] $\beta_i(k) = \frac{\epsilon_i}{b+\sum_{j=1}^{m_k}\epsilon_j}$ and $\beta_0(k) = \frac{b}{b+\sum_{j=1}^{m_k}\epsilon_j}$, where $\epsilon_i \triangleq \exp[-\frac{\nu_i'(k)S^{-1}(k)\nu_i(k)}{2}]$ and $b \triangleq (2\pi/\gamma)^{n_z/2}\lambda V_k c_{n_z}(1 - P_D P_G)/P_D$. For our scenarios, where there is exactly one measurement at each time step, $m_k = 1$, and the PDA tries to determine whether the measurement is from the target of interest or not.

D. Kolmogorov-Smirnov Tests

Kolmogorov-Smirnov (KS) tests provide another approach for determining the probability of whether a measurement in a set originated from the object of interest [7], [8]. Consider independent observations $\bar{x}_1, \bar{x}_2, ..., \bar{x}_n$ of a random variable with unknown cumulative distribution function (CDF) $F(\bar{x})$. If the null hypothesis is $H_0 : F(\bar{x}) = F_0(\bar{x})$, then any test of this hypothesis is a goodness-of-fit test [1]. The KS tests and many simple variants are goodness of fit tests. Empirical CDFs are formed for a "window" of n samples and some metric is then applied to measure the distance between theoretical and empirical CDFs.

For n samples ordered such that $\bar{x}_{(1)} \leq \bar{x}_{(2)} \leq \cdots \leq \bar{x}_{(n)}$, the empirical cumulative distribution function is [1]

$$S_n(\bar{x}) = \begin{cases} 0, & \bar{x} < \bar{x}_{(1)} \\ r/n, & \bar{x}_{(r)} \leq \bar{x} < \bar{x}_{(r+1)} \\ 1, & \bar{x}_{(n)} \leq \bar{x} \end{cases}. \tag{6}$$

If $F_0(\bar{x})$ is the true, fully specified theoretical CDF from which the samples are drawn, then from the strong law of large numbers $\lim_{n\to\infty}\{S_n(\bar{x}) = F_0(\bar{x})\} = 1$.

Define the following metric for measuring the separation between theoretical and empirical CDFs:

$$A_n = \left| \int_0^{\infty} (F_0(\bar{x}) - S_n(\bar{x}))d\bar{x} \right|. \tag{7}$$

This is different from the usual measure of deviation

$$D_n = \sup_{\bar{x}} |S_n(\bar{x}) - F_0(\bar{x})| \tag{8}$$

and its variants that are well developed for the KS test. The measure of deviation (8) proved to not be sensitive enough for the data association problem considered here. The metric (7) is more sensitive to incorrect measurements. Equation (7) largely takes advantage of how a single incorrect measurement affects the shape of the empirical CDF.

There has been a large body of theory developed using D_n in (8), demonstrating many properties, such as the distribution of D_n is independent of the distribution $F_0(\bar{x})$ for continuous CDFs [3]. Critical values or thresholds $d_n(\alpha)$ have also been established [1], [5], [6] such that $P\{D_n < d_n(\alpha)\} = 1 - \alpha$, where $1 - \alpha$ is the confidence level. That is, if $D_n < d_n(\alpha)$, then with probability $1-\alpha$, the empirical CDF $S_n(\bar{x})$ is formed from n samples drawn from $F_0(\bar{x})$. There is currently no theoretical basis for specifying critical values for A_n in (7), and this is an area of future work.

In this test, a CDF is formed using a "window" of n samples, $\bar{x}_i, \bar{x}_{i+1}, ..., \bar{x}_{i+n-1}$. This CDF is compared to an empirically formed "theoretical" CDF. The value of A_n is computed from this comparison. If the value of A_n is above a preset threshold, it is assumed that the current window of n samples most likely contains at least one due to an incorrect measurement. This process is then repeated for the next window of n samples, i.e., $\bar{x}_{i+1}, \bar{x}_{i+2}, ..., \bar{x}_{i+n}$. If the value of A_n is above the threshold for n consecutive shifts of the window of samples, the location of an incorrect measurement is uniquely determined (assuming a uniform distribution of incorrect measurements where each window contains a single incorrect measurement). After determining the location of incorrect measurements, these measurements are removed from the list of measurements and the remaining measurements are used in the filtering algorithm, taking into account the time signatures associated with each measurement. Section V will show the performance of using (7) on the example discussed next.

IV. Example

Consider two cars traveling in a small convoy. We use only two cars in this example for simplicity and to more clearly illustrate the performance differences of the different algorithms. Extension to more vehicles is easily accomplished because we are assuming that the sensor is attempting to follow the car of interest as best it can. This implies that the percentage of incorrect measurements made should only be a function of the average separation distance between the car of interest and those nearby and not of the total number of cars in the convoy. The effect of separation distance between the two vehicles will be explored. Note that the measurement noise in the system is a characteristic of the sensor and is thus independent of the number of cars in the convoy.

The dynamics of this system are modeled using constrained Euler-Lagrange equations [9]. To obtain the Euler-Lagrange equations, start by defining the *Lagrangian* for the generalized state q, defined as the kinetic energy T minus the potential energy V: $L(q, \dot{q}) = T(q, \dot{q}) - V(q)$.

The *Euler-Lagrange equations* in vector form are

$$\frac{d}{dt}\frac{\partial L}{\partial \dot{q}} - \frac{\partial L}{\partial q} = \rho \omega^T(q) \tag{9}$$

where ρ is the *Lagrange multiplier* and $\omega(q)$ is the constraint matrix. We make use of the *Lagrange-d'Alembert formulation*, i.e. the forces of constraint do no virtual work, to write down the *Pfaffian constraint* [9] equation $\omega(q)\dot{q} = 0$.

Since we assume the cars to be fixed to a flat road, thus the only terms in the Lagrangian are those associated with the kinetic energy. This means that the cars do not have any Z (vertical) component to their motion, i.e., $q = (X, Y)$. Thus, $L(X, Y, \dot{X}, \dot{Y}) = M(\dot{X}^2 + \dot{Y}^2)/2$, where M is the mass of the car and X and Y are the Cartesian coordinates.

The model $g(x)$ of the road is assumed, yielding the constraint equation $y(x) = g(x)$. We obtain $\omega(q)$ by differentiating the constraint equation with respect to time and writing the result in the Pfaffian constraint form. The Lagrange multiplier is found by simultaneously solving (9) along with $\frac{d^2}{dt^2}(y(x(t)) = g(x(t)))$.

For our example, the state $x(t) = [X(t), Y(t)]$ and $g(x) = \sin(X)$. There is no physical significance to this choice of constraint; any analytical function of the configuration variables could be used in place of $\sin(X)$. Choosing $g(x) = \sin(X)$, ω and ρ are $\omega(x(t)) = \{-\cos(X(t)), 1\}$ and $\rho = -\frac{M\sin(X(t))X'(t)^2}{1+\cos(X(t)}$. The resulting dynamic equations are $\ddot{X}(t) = \frac{\cos(X(t))\sin(X(t))X'(t)^2}{1+\cos(X(t))^2}$ and $\ddot{Y}(t) = -\frac{\sin(X(t))X'(t)^2}{1+\cos(X(t))^2}$. Thus the complete continuous-time equations of motion for the system are defined. The truth trajectories are obtained by integrating these continuous equations (with noise). In practice it is assumed that we do not know the true position values, but in simulation, we use these values for creating "measurements" and for analysis. The discrete model (1) is obtained by taking the discrete approximation of the continuous dynamics using the time step $T = 0.0001$s.

For this example, the process noises are drawn from the normal distributions $\mathcal{N}$[0 m/s, 0.1 m/s] and $\mathcal{N}$[0 m, 0.1 m]. Measurements of X and Y are generated assuming the measurement error is normally distributed $\mathcal{N}$[0 m, 0.5 m]. A single set of measurements is then created by randomly inserting with probability P_D at each time step the measurement due to the car of interest and inserting with probability $1 - P_D$ the measurement from the "other" car. $P_D = 0.9$ for the results shown in Section V. Once the mixed list of measurements is obtained, gating and the PDA are relatively straightforward processes. The actual implementation for our specific case will be further addressed in the next section.

For the KS test, the distribution of the magnitude of the innovations is used. A theoretical CDF is determined by generating a large number ($n = 20,000$) of true measurements as explained above, then computing the magnitude of the innovations and from this generating the theoretical CDF. In practice, "theoretical" CDFs can also be generated "empirically" in a similar manner when the object of interest is known not to be near any other objects. In this way, the KS test does not depend on knowledge of the error distributions.

V. Results

All measurements for the two-car system are Cartesian X and Y measurements, as stated in the problem definition. All the results in this section will be stated in a "radial" format, where each distance presented is actually a Euclidean norm. All RMS errors represent estimation errors.

A. Gating

The γ parameter in (4) sets the gate size. Our goal is to exclude the false measurements while keeping as many of the true measurements as possible. We set $\gamma = 9$ thus allowing about 98.9% of the true measurements into the gate [2]. While the value of γ has been chosen heuristically, optimum methods for calculating this threshold do exist [4].

Figure 2 shows gating performance as a function of the separation distance between the two vehicles, for the given gate size. The performance is measured by the percentage of

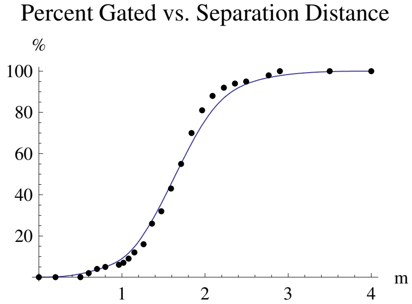

Fig. 2. Percent of incorrect measurements correctly gated out. The points are data points obtained by varying the separation distance between the two cars; the solid line is a curve fit to the data.

the incorrect measurements that are correctly gated out. For example, looking at Figure 2, under about 0.5 m of separation between vehicles, none of the incorrect measurements are gated out; all 10% of the incorrect measurements will be left in the measurement list. In the same way, above about 3.5 m, all of the incorrect measurements are gated out. Thus, with a standard deviation for the measurements of $\sigma = 0.5$ m, at around seven σ's separation (3.5 m), gating works perfectly.

B. Kolmogorov-Smirnov tests

The objective of the variation of the Kolmogorov-Smirnov test described in Section III is to isolate measurements not coming from the car of interest. As described, this is accomplished by calculating the integral of the difference between theoretical and empirical CDFs as in (7), and comparing this value against a threshold.

Limited knowledge of the sensor characteristics is assumed. We have 20,000 samples taken from the car we of interest during a period when it is known that no other cars are on the road. This stream of measurements is used to form the theoretical CDF.

The window size for the two-car example is set to ten, i.e., $n = 10$. Thus, an empirical CDF function is formed according to (6) for the first ten measurements from the sensor. This function is used to obtain a value of A_n according to (7). This value is then compared against a threshold, set by inspection. Results from this process iterated over the first 210 measurements are shown in Figure 3. The incorrect measurements are at positions 28, 59, 89, 117, 146, and 178. In Figure 3 we can see "humps" in the values of A_n centered around these positions. This happens because there are actually 10 windows that will span the incorrect measurement. When the cars are far from each other, at 7 m apart, these humps are very pronounced and are well above the threshold. When the cars are closer together, at 2 m, the incorrect measurements do not produce values of A_n that clearly exceed the threshold.

Applying the standard KS test metric (8) for this example does not isolate single measurements. This can be seen in Figure 4, where there is no clear difference between a window that contains an incorrect measurement and one that does not, i.e., there are no "humps." Thus, setting a clear threshold is not possible.

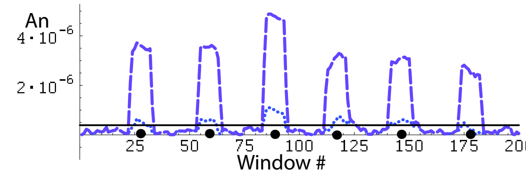

Fig. 3. Integral difference (An) between empirical and theoretical CDF's as a function of window number. The horizontal axis is the window number, where 1 corresponds to the window covering measurements 1 through 10, 2 corresponds to measurements 2 through 11, and so on. The vertical axis is the value of A_n for that window. The solid line is the threshold. The dashed line represents when the two cars are 7 m apart. The dotted line represents the situation when the two cars are 2 m apart. The solid dots represent the positions of incorrect measurements.

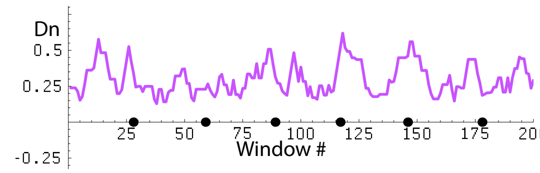

Fig. 4. Maximum difference (Dn) between empirical and theoretical CDF's as a function of window number. The solid line is the maximum difference between theoretical and empirical CDFs over a window of 10 measurements. The solid dots represent the positions of incorrect measurements.

Future work for this test includes applying more advanced filtering techniques such that the "humps" in the variant of the KS test are more pronounced at lower separation distances. The threshold could then be lowered without affecting the truth data and better performance can be achieved.

C. PDA

PDA results are presented using PDA alone, with the KS test variant, and with gating. As mentioned previously, the PDA for each case will need to be initialized in the same way as the EKF. P_D and P_G need to be set for the PDA. As before, $P_D = 0.90$. For the case of PDA alone or PDA used with the KS variant, $P_G = 1$. When there is gating, $P_G = 0.989$ (because $\gamma = 9$ in the gating, see [2]).

When the cars are far apart, the PDA will have lower RMS error when used in conjunction with gating or the KS test as opposed to using PDA alone. In Figure 5, after settling PDA with the KS test has lower RMS error than PDA with gating. This occurs because at a distance of 10 m, the KS

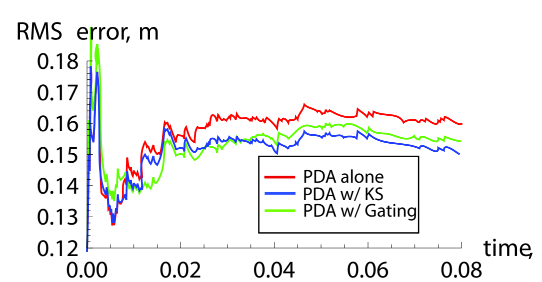

Fig. 5. RMS errors for PDA with the KS test, PDA with gating, and PDA alone when two cars are moving 10m apart.

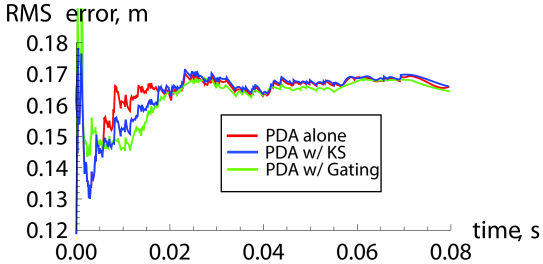

Fig. 6. RMS errors for PDA with the KS test, PDA with gating, and PDA alone when two cars are moving 2m apart.

test removes all of the incorrect measurements while leaving the remaining true measurements unaffected. Gating will also remove all of the incorrect measurements at this distance, but will remove a percentage of the remaining true measurements (the percentage is a function of the gate size, which for this case is 1.1%).

Results for a separation distance of 2 m are shown in Figure 6. The RMS errors of the PDA alone, PDA with the KS test, and PDA with gating are much closer. This can be explained by looking back at Figures 2 and 3. At 2 m of separation between vehicles, gating will not "gate out" all of the incorrect measurements. In fact, for this separation distance, about 20% of the incorrect samples will still be left in the list of measurements, which will cause the RMS error to be higher than in cases where 100% of the incorrect measurements are correctly identified. At 2 m of separation between vehicles, the KS test variant has difficulty in differentiating between correct and incorrect measurements, i.e., it is difficult to set a reliable threshold. This difficulty leads to a number of correct measurements being falsely identified as incorrect, leading to higher RMS error after filtering. In terms of RMS error after settling, PDA with gating does marginally better than PDA with the KS test due to a higher percentage of true measurements being present in the former.

VI. Conclusions & Future Work

We have shown that the data association ideas of gating, the probabilistic data association filter, and a Kolmogorov-Smirnov type test provide an effective manner in which to deal with the problem of ambiguous measurements. It has been shown that in the situation of mixed measurements, these methods can be applied to provide noticeably improved tracking over the information provided by the measurements alone.

The effectiveness of these data association methods applied to this problem suggest several avenues of future work. The first being to formally address the problem from a more geometrical interpretation. Although we have been considering a nonlinear system, we have been assuming that the various noise distributions vary nicely along the trajectories. In the two car example, it was assumed that the road was $\Re^2$, not taking into account the underlying manifold structure where the cars are constrained to the road. In the future, we would like to show that improved tracking performance is possible when the distributions are dealt with in a more geometrically correct way.

References

[1] K. Ord A. Stuart and S. Arnold. *Kendall's Advanced Theory of Statistics*. Oxford University Press, 1999.

[2] Y. Bar-Shalom and T.E. Fortmann. *Tracking and Data Association*. Academic Press, Inc., 1988.

[3] W. Feller. On the kolmogorov-smirnov limit theorems for empirical distributions. *Annals Math. Stat.*, 19:177–189, 1948.

[4] Y. Kosuge and T. Matsuzaki. The optimum gate shape and threshold for target tracking. In *SICE 2003 Annual Conference*, Fukui, Japan, 2003.

[5] F.J. Massey. A note on the power of a non-parametric test. *Annals Math. Stat.*, 21:440–443, 1950.

[6] L.H. Miller. Table of percentage points of kolmogorov statistics. *J. Amer. Stat. Assoc.*, 51, 1956.

[7] R. M. Powers and L. Y. Pao. Using kolmogorov-smirnov tests to detect track-loss in the absence of truth data. *Proc. IEEE Conf. Decision and Control and European Control Conf.*, pages 3097–3104, 2005.

[8] R. M. Powers and L. Y. Pao. Power and robustness of a track-loss detector based on kolmogorov-smirnov tests. *Proc. American Control Conf.*, pages 3757–3764, 2006.

[9] Z. Li R.M. Murray and S. Shankar Sastry. *A Mathematical Introduction to Robotic Manipulation*. CRC Press, 1994.

[10] H. Durrant-Whyte T. Bailey, B. Upcroft. Validation gating for non-linear non-gaussian target tracking. In *Ninth International Conference on Information Fusion*, Florence, Italy, 2006.