

A Particle Filtering Approach to FM-Band Passive Radar Tracking and Automatic Target Recognition

Shawn Herman and Pierre Moulin

Department of Electrical and Computer Engineering
Beckman Institute, University of Illinois at Urbana-Champaign
405 N. Mathews
Urbana, IL 61801, USA
217-244-1089
{smherman,moulin}@ifp.uiuc.edu

Abstract—We present two stochastic filters for an FM-band passive air surveillance radar. The first system uses an extended Kalman filter and delay-Doppler measurements to track targets. The second system uses a particle filter to simultaneously track *and* classify targets. Automatic target recognition is made possible by the inclusion of radar cross section (RCS) in the measurement vector. The extended Kalman filter cannot take advantage of radar cross section measurements because there is no closed-form relationship between the state elements which determine target aspect and the resulting RCS measurement. We believe that this is the first work to propose the use of RCS for the purpose of target recognition within a passive radar system. We also present many simulation results for a challenging 2-target 3-sensor task involving trajectories which nearly coincide for a portion of their track length.

TABLE OF CONTENTS

- 1 INTRODUCTION
- 2 BACKGROUND IN ELECTROMAGNETICS
- 3 KALMAN FILTER-BASED TRACKING
- 4 PARTICLE FILTER-BASED TRACKING AND CLASSIFICATION
- 5 EXPERIMENTAL RESULTS
- 6 CONCLUSIONS
- 7 ACKNOWLEDGEMENTS

1. INTRODUCTION

Radar systems extract information pertaining to position and velocity by probing targets within their surveillance volume using electromagnetic (EM) waves. The range of frequencies of the EM waves employed by such systems can vary considerably. A passive radar might utilize signals in the FM band (100 MHz), while a typical active radar could operate at 10 GHz (X-band). Further still, a radar with infrared sensors might use waves with frequencies in excess of 10^{13} Hz. In all cases, information about a target

is obtained by measuring various aspects of the EM waves which have scattered off the target and traveled back to the radar receiver. For example, time delay between transmission and reception of an EM pulse can provide an estimate of the range to the target, while Doppler shift in carrier frequency can be related to target velocity. In addition, the power of the scattered wave, as compared to that of the transmitted wave, can provide an estimate of the electromagnetic “size” of the target (i.e., the effective electromagnetic capture area of the target).

In this paper, we are interested in exploring the use of commercial FM broadcasts for tracking and classification of airborne targets. While delay and Doppler shift of a scattered FM waveform can be used to estimate position and velocity, these measurements do not convey much in the way of target-specific information. Thus, while it is true that an F-16 can execute certain maneuvers that a DC-10 cannot, in general, the kinematic information supplied by delay and Doppler is insufficient to identify an unknown target. What is needed is another type of measurement. In this paper, we propose the use of radar cross section (RCS) as this additional measurement. Our interest in FM-band radar is due to the resulting frequency of operation. At FM-band frequencies, RCS varies much slower with target aspect angle than at X-band, and this provides a more robust feature for classification. Our task then is to incorporate RCS as an additional component in our measurement vector. Unfortunately, we will show in Section 2 that this is not practical with a Kalman filter (or an extended Kalman filter, for that matter). Instead, we resort to the sample-based nonparametric density estimation technique known as particle filtering [1, 2]. In addition to permitting the use of RCS as a component of our measurement vector, we will show that particle filtering also presents a useful framework for a joint approach to tracking and classification.

A particle filtering approach could also be used for joint tracking and classification in other operational scenarios, such as television-based passive radar or X-band active radar. However, depending on the characteristics of the illuminating waveform, a set of measurements other than delay, Doppler shift, and RCS might be either preferable or

necessary. The success of any joint tracking/classification system, including the FM-band passive radar presented in this paper, relies upon the fidelity with which the measurement set conveys information regarding target location and identity. Thus, while the specific design equations for our FM-band passive radar (which assume measurements of delay, Doppler, and RCS) might need to be altered, the overall design principle for a particle filtering approach to joint tracking/classification would remain unchanged.

This work was influenced by that of Miller *et al.* [3, 4] who also formulated a joint approach to tracking and classification but used jump-diffusion processes as their inference engine. In addition, many papers have been published in the last decade which explore the use of particle filters for various tracking applications [5–9]. We believe our work is significant, though, for two reasons. First, we believe we are the first to publish a formulation for a tracking filter which relies upon commercial FM broadcasts.¹ Second, and more importantly, we believe that we are the first to design a particle filter which uses radar cross section as a feature for automatic target recognition.

The rest of this paper is organized as follows. Section 2 presents relevant aspects of scattering theory including the definition of radar cross section. It highlights the complex relationship between target aspect and RCS. In Section 3, we present an FM-band air surveillance radar based upon the extended Kalman filter. This system will be the benchmark for comparison with the particle filter. We will discuss in detail its system and measurement models, as well as its method of data association. In Section 4, we provide a similar development for the particle filter, paying particular attention to the problem of data association and the incorporation of RCS into the measurement vector. Next, experimental results obtained through Monte Carlo simulation of a 2-target 3-sensor scenario are presented in Section 5. In addition to showing that RCS-based classification is feasible, these results also demonstrate that RCS measurements can improve tracking performance by alleviating ambiguity during data association. Finally, we conclude in Section 6.

2. BACKGROUND IN ELECTROMAGNETICS

The interaction between target and waveform in radar systems is properly addressed using electromagnetic scattering theory. In this section, we present an overview of the EM concepts pertinent to our application. As a means of motivating this discussion, we begin by introducing the operational scenario we will address: ground-based FM-band passive air surveillance. A discussion of FM-band passive radar then leads naturally to consideration of radar cross section (or, more precisely, received power) as a potential feature for classification.

¹In [10], Howland presents a tracking formulation which uses a single television broadcast to provide measurements of target bearing and Doppler shift.

Ground-Based FM-Band Passive Radar

The majority of modern radar systems are *active*. An active radar both transmits and receives electromagnetic waveforms for the purpose of detection and estimation. In contrast, a *passive* radar system does not transmit electromagnetic energy of its own; instead, it relies on “illuminators of opportunity” such as commercial FM radio or television broadcasts. Regardless of the source and nature of EM transmission, both types of radar attempt to measure the scattering of these waveforms in order to detect the presence of targets within their surveillance volume. Individual detections (in measurement space) which are determined to have originated from the same target are associated and referred to as a *track*. Measured data corresponding to a given track can then be stochastically filtered to produce an estimate of the unknown flight path.

Although most radar systems are active, passive radar can provide some advantages, especially in military contexts [11–13]. Because an active radar advertises its location through its transmissions, it can quickly become a combat target. A passive radar, on the other hand, offers the primary benefits of survivability and robustness against deliberate directional interference. Second, since commercial transmitters focus their broadcast energy toward the Earth’s surface, a passive radar permits illumination of low-flying aircraft. Finally, the signals used by passive radar lie in the FM, VHF, and UHF frequency ranges (55–885 MHz). Because the wavelengths at these frequencies are orders of magnitude larger than airborne precipitation, these signals are virtually immune to weather-induced degradation [11].

The main disadvantage of passive radar is a matter of signal design. As mentioned, a passive radar gathers information about targets using commercial broadcasts, most of which are far from optimal in terms of tracking capabilities. To understand why, consider that an active radar estimates range by measuring the time delay between pulse transmission and reception of the target echo. In contrast, commercial broadcasts are not pulsed waveforms; they are continuous signals which are always “on”. In the case of a passive radar using FM broadcasts, target range must be estimated by correlating the modulation waveform of the direct FM signal with the modulation waveform of the scattered signal. As such, the accuracy of this range measurement, or *range resolution*, is related to the maximum excursion of the frequency modulation. More generally, it can be shown that the achievable range resolution is inversely proportional to the bandwidth of the illuminating signal [14]. Since a chirp waveform used by an active radar can have a bandwidth which is 10,000 times larger than that of an FM radio broadcast, the challenge inherent in FM-band passive radar is apparent.

Having established that the tracking performance of passive radar may be substantially worse than that of its active counterpart, a logical question might be, “Are there any ad-

vantages to using passive radar for non-military applications?” As it turns out, the FM frequency band is well-suited to target classification. Before we can explain why this is so, we first must introduce the concept of radar cross section.

Radar Cross Section

In radar applications, we are often interested in how well a target captures and reradiates electromagnetic energy. Radar cross section (RCS) is a measure of the magnitude of this reflection process expressed as an effective area. RCS is a function of the size, shape, and composition of the target as well as the polarization and frequency of the incident wave. RCS also depends on the position and orientation of the target with respect to both the source of the incident wave (e.g., the FM transmitter) and the measurement location (i.e., the radar receiver). RCS is formally defined as

$$\sigma(\mathbf{k}_s, \mathbf{k}_i) = \lim_{R \rightarrow \infty} 4\pi R^2 \frac{|\mathbf{E}^s(\mathbf{k}_s)|^2}{|\mathbf{E}^i(\mathbf{k}_i)|^2}, \quad (1)$$

where $|\mathbf{E}^i|$ is the amplitude of the incident wave measured at the target and $|\mathbf{E}^s|$ is the amplitude of the scattered wave measured at a distance R from the target [15]. \mathbf{k}_i and \mathbf{k}_s denote the incident and scattered directions, respectively, as measured in a coordinate system which is centered on the target. The limiting process assures that the scattered wave is measured in the far-field.² The incident wave in the RCS calculation is commonly taken to be a plane wave. This assumption is justified when the distance between the transmitter and the target is large because spherical phase fronts can then be approximated as locally planar.

Since RCS is a function of size, shape, and composition of the target, it is a potential feature for classification. However, a radar cannot measure RCS directly. Instead, it measures the power of the received waveform P_r which is related to RCS through the radar range equation

$$\sigma = \frac{(4\pi)^3 R_t^2 R_r^2}{\beta_t P_t \beta_r \lambda^2} P_r, \quad (2)$$

where P_t is the transmitted power, λ is the wavelength, R_t is the transmitter-to-target range, and R_r is the receiver-to-target range.³ β_t and β_r are antenna gains for the transmitter and receiver, respectively. The fraction in (2) can be viewed as the normalization term needed to remove the effects of transmission and propagation from the received waveform, leaving a quantity (RCS) which depends on the target alone. It is reasonable to assume that P_t , β_t , and β_r are either known quantities or may be determined through suitable calibration. The two distances, R_t and R_r , must be approximated using an estimate of the target’s position (which is available, of course, in a joint tracking/classification framework).

²For notational simplicity, we did not indicate the dependence of RCS on frequency and polarization in (1).

³The wavelength used in (2) corresponds to the carrier frequency of the broadcast, since that is where most of the transmitted power is concentrated.

Now, we return to our claim that FM radio frequencies are well-suited for classification. As we have just seen, the RCS of a target will vary as its position and orientation with respect to the transmitter or receiver changes. To ensure robust classification in the presence of noise, which may induce errors in our estimates of position and orientation, it is helpful that RCS vary “slowly” with small changes in these components of the target’s state. The variation in RCS, as reflected by the number of nulls encountered as a target’s aspect changes, is proportional to the electrical length of the target. At FM-band frequencies (100 MHz), a fighter-sized aircraft is approximately five wavelengths long. In contrast, at the X-band frequencies used by active radars (10 GHz), the same aircraft would be 500 wavelengths long! Therefore, although FM-band frequencies yield poor range resolution, they have the potential for robust classification through measurement of RCS [16, 17].

Method of Moments

We have just shown that radar cross section may be a viable feature for classification within an FM-band passive radar system. However, in order to use RCS (as measured by received power), we will need to know what the expected RCS is for each target in our library, over the entire range of frequencies and aspect angles which may be encountered during classification. According to (1), this amounts to solving for the scattered far-field $\mathbf{E}^s(\mathbf{k}_s)$, given the incident frequency and $\mathbf{E}^i(\mathbf{k}_i)$. To accomplish this, the current induced by the incident wave on the surface of the target, \mathbf{J}_s , must be found. Unfortunately, for all but the simplest target geometries, it is impossible to find an analytic expression for \mathbf{J}_s . Instead, we must rely on either empirical or numeric techniques. In this section, we will discuss one of the most popular approaches to numerically estimating the RCS of a target. Our goal is to highlight the complicated relationship that exists between target state (e.g., position and orientation) and RCS.

The method of moments (MoM) is a numeric technique which has found widespread use in the solution of scattering problems involving complex targets [18]. The method of moments applies to general linear-operator equations, such as

$$\mathcal{L}r = e, \quad (3)$$

where \mathcal{L} is a linear operator, r is an unknown response, and e is a known excitation. The unknown response r is expanded as the sum of basis functions $r = \sum_{j=1}^N a_j r_j$. To solve for the N unknowns $\{a_1, \dots, a_N\}$, we write N linearly independent equations. These are obtained by taking the inner product of (3) with a set of N testing functions $\{y_1, \dots, y_N\}$,

$$\langle y_i, \mathcal{L}r \rangle = \sum_{j=1}^N a_j \langle y_i, \mathcal{L}r_j \rangle = \langle y_i, e \rangle, \quad i = 1, \dots, N. \quad (4)$$

This equation can be written in matrix notation as $Z \mathbf{a} = \mathbf{e}$

where $Z_{ij} = \langle y_i, \mathcal{L}r_j \rangle$, \mathbf{a} is a column vector of the unknown coefficients, and $e_i = \langle y_i, e \rangle$. Matrix inversion of Z yields the desired basis coefficients: $\mathbf{a} = Z^{-1}\mathbf{e}$.

In our electromagnetic scattering application, \mathcal{L} is an integrodifferential operator derived from one of the boundary conditions of Maxwell's equations, e is an incident field term (either electric or magnetic), and r is an approximation of \mathbf{J}_s . The shape of the target is provided to the MoM code by a CAD file which partitions the surface of the target into smooth patches or planar facets. Both the basis functions $\{r_i\}$ and the testing functions $\{y_i\}$ are then defined in terms of the facetization of the target. A popular choice of basis is defined on pairs of adjacent triangular facets [19]. The so-called RWG basis is designed to maintain continuity of the normal component of \mathbf{J}_s across facet boundaries, thereby avoiding the need for fictitious line and point charges. A popular choice of testing function is $y_i = r_i$, commonly referred to as Galerkin's method. If a target is modeled as a closed surface, use of the RWG basis requires one function per triangular facet edge, yielding a number of unknowns proportional to the number of facets in the CAD representation. Since the level of facetization is dictated by the electrical length of the target, an accurate representation of the surface current on a fighter-sized aircraft at gigahertz frequencies may require millions of facets.⁴ Unfortunately, with this many unknowns, it is clearly impractical to invert the matrix Z . In cases such as these, a conjugate gradient technique can be used to replace the $O(N^3)$ matrix inversion with an $O(KN^2)$ iterative approach requiring only matrix-vector multiplies. (K is the number of iterations required for convergence.) If the amount of computation is still unwieldy, the Multilevel Fast-Multipole Algorithm (MLFMA) can be used to reduce the number of operations in the matrix-vector multiply, resulting in $O(KN \log N)$ computation [20].

For the experimental results we present in Section 5, all RCS values were computed numerically using an MLFMA-based software package called FISC (Fast Illinois Solver Code) [21]. FISC finds \mathbf{J}_s as described above and then solves for the scattered electric field using the far-field radiation integral

$$\mathbf{E}^s(\mathbf{k}_s) = -j\omega\mu \frac{e^{-jkR}}{4\pi R} \oint_S (\mathbf{J}_s - (\hat{n} \cdot \mathbf{J}_s)\hat{n}) e^{j\mathbf{k}_s \cdot \mathbf{r}} dS, \quad (5)$$

where μ is the permeability of the medium, \mathbf{r} is the location in a target-centered coordinate system of an infinitesimal patch on the surface of the scatterer, and \hat{n} is the outward surface normal at \mathbf{r} . The wave number k is just the magnitude of \mathbf{k}_s , $k = \frac{2\pi}{\lambda}$. It is beneficial to pause at this point in our discussion and draw attention to the fact that there is no simple closed-form relationship between a target's aspect and the resulting RCS. Thus, even though an extended Kalman filter can handle nonlinear relationships between

⁴As a rule of thumb, the maximum length of a facet edge should be between 0.1 and 0.2 wavelengths.

state (target aspect) and measurement (RCS), that relationship must still be made explicit. In the case of RCS, the closest we could come would involve storing the entire admittance matrix (Z^{-1}) and manipulating a discretized version of (5), an approach which is clearly impractical for real-time operation.⁵

Up until now, we have overlooked polarization in our discussion. However, it must be taken into account in order to understand how FISC calculates RCS values. Recall that \mathbf{E}^i specifies the amplitude, phase, and orientation of the incident electric field at each point in space. Since RCS is defined as a ratio of magnitudes and \mathbf{E}^s is related to \mathbf{E}^i through integrodifferential equations, the amplitude and phase of the incident field are arbitrary. As such, FISC sets $|\mathbf{E}^i| = 1$. The only parameter of the incident field which matters then is its orientation (and frequency). If we assume in the RCS computation (as FISC does) that \mathbf{E}^i is a uniform plane wave traveling through free space, the phase fronts are planes perpendicular to the direction of propagation, $\mathbf{E}^i \cdot \mathbf{k}_i = 0$. As such, \mathbf{E}^i can be decomposed into two orthogonal components

$$\mathbf{E}^i = E_\theta^i \hat{i}_\theta + E_\phi^i \hat{i}_\phi, \quad (6)$$

where (θ, ϕ) represent the incident direction in a spherical coordinate system centered on the target. Because the system relating \mathbf{E}^i and \mathbf{J}_s is linear, FISC only needs to compute \mathbf{J}_s for $\mathbf{E}^i = \hat{i}_\theta$ and $\mathbf{E}^i = \hat{i}_\phi$; the current distribution due to all other incident orientations can be synthesized from these two results. \hat{i}_θ and \hat{i}_ϕ are typically referred to as "vertical" and "horizontal" polarizations, respectively.

In an analogous fashion, the scattered electric field can be expressed as a sum of orthogonal components

$$\mathbf{E}^s = E_{\theta'}^s \hat{i}_{\theta'} + E_{\phi'}^s \hat{i}_{\phi'}, \quad (7)$$

where the spherical angles (θ', ϕ') in (7) are primed because they correspond to the scattered direction, which is not necessarily the same as the incident direction.⁶ We can see that given the incident and scattered directions, there are only four coefficients which are needed to calculate \mathbf{E}^s for any \mathbf{E}^i satisfying $\mathbf{E}^i \cdot \mathbf{k}_i = 0$. These four quantities constitute entries in a scattering matrix S ,

$$\begin{bmatrix} E_{\theta'}^s \\ E_{\phi'}^s \end{bmatrix} = \underbrace{\begin{bmatrix} s_{vv} & s_{vh} \\ s_{hv} & s_{hh} \end{bmatrix}}_S \begin{bmatrix} E_\theta^i \\ E_\phi^i \end{bmatrix}, \quad (8)$$

where we have adopted the popular notation of using 'v' and 'h' to denote vertical and horizontal polarizations.

⁵If a 3-dimensional bistatic scattering center model could be found which was accurate over a wide range of target aspects, its gradient could be used in the measurement update for the extended Kalman filter. However, no such model currently exists.

⁶If $\mathbf{k}_s = -\mathbf{k}_i$, the spherical angles *will* be the same and the scenario is termed monostatic scattering. The general case is referred to as bistatic scattering.




Figure 1. RCS plots for two targets at 99.9 MHz and zero degrees elevation. The x -axis in each plot corresponds to incident azimuth angle while the y -axis corresponds to scattered azimuth angle. Left plot: Falcon-20 commercial jet. Right plot: VFY-218 fighter.




Figure 2. (a) Example of a positive elevation angle (EL) on the Cartesian axes. (b) Example of a positive azimuth angle (AZ) on the Cartesian axes.

Since each element of S varies with \mathbf{k}_i and \mathbf{k}_s , the scattering matrix is actually a collection of four complex-valued functions. The square magnitude of each coefficient yields the corresponding RCS (e.g., $\sigma_{vh} = |s_{vh}|^2$). It is specifically these scattering coefficients that FISC calculates. RCS plots for two different aircraft are shown in Fig. 1. Note that in order to plot RCS in two dimensions, we have set both incident and scattered elevation angles to zero degrees. The plots depict σ_{hh} , the s_{hh} polarization component of RCS.

RCS Databases and Sampling Considerations

In the previous section, we discussed how a software package such as FISC can be used to numerically approximate the scattering matrix for a complex target such as an aircraft. For each class within our library of potential targets, we will use FISC to generate a database of RCS values (offline, prior to tracking) corresponding to various incident and scattered directions. We can then estimate RCS along an arbitrary flight path by interpolating between database entries.

FISC specifies the direction of wave propagation using elevation and azimuth angles (EL, AZ) in a target-centered coordinate system. These angles are related to the spherical

coordinate system by $EL = 90^\circ - \theta$ and $AZ = -\phi$. An example of positive elevation and azimuth angles is provided in Fig. 2. The incident wave's direction can then be specified by the pair (EL^i, AZ^i) while the scattered wave direction is indicated by (EL^s, AZ^s) . When computing RCS values for more than one pair of incident and scattered directions, FISC imposes the requirement that all angular values be separated by multiples of fixed elevation and azimuth increments. Define Δ_{EL} as the fixed elevation increment, Δ_{AZ} as the fixed azimuth increment (where $\Delta_{EL} \neq \Delta_{AZ}$ necessarily), and Δ_f as the frequency increment. In this section, we will consider how to choose these three values, with the intent of minimizing the size of each database by using the largest increments possible.

To gain insight into this choice of increments, we will assume that our target can be modeled as a collection of ideal scattering centers, each with value g_n and position $\mathbf{r}_n = (x_n, y_n, z_n)$ [22]. This model is appropriate when the wavelength of the incident field is small relative to the target dimensions. If we also assume that each scattering center is visible from all observation angles, the monostatic scattering from the target can be written as

$$G(f, \theta, \phi) = \sum_n g_n \exp[-j2\mathbf{k} \cdot \mathbf{r}_n]. \quad (9)$$

If we substitute the following expression for the wave vector

$$\mathbf{k} = \frac{2\pi}{\lambda} \left(\sin \theta \cos \phi \hat{i}_x + \sin \theta \sin \phi \hat{i}_y + \cos \theta \hat{i}_z \right) \quad (10)$$

into (9), we arrive at

$$G(f, \theta, \phi) = \sum_n g_n \exp[-j2\pi(Xx_n + Yy_n + Zz_n)], \quad (11)$$

$$\text{where } X = \frac{2f}{c} \sin \theta \cos \phi, \quad Y = \frac{2f}{c} \sin \theta \sin \phi, \\ Z = \frac{2f}{c} \cos \theta, \quad (12)$$

and c is the speed of light. The capitalized notation highlights the fact that (x, y, z) and (X, Y, Z) form a Fourier transform pair. Since building an RCS database amounts to uniform sampling of $G(f, \theta, \phi)$, a reasonable choice for stepsizes would be ones which satisfy the Nyquist criterion for G .⁷

To enforce the Nyquist criterion, we need to make some assumptions about the spatial extent of our target as represented by the scattering centers $\{g_n\}$. We assume that our target is oriented so that its nose points along the positive x -axis and the left wing (from the pilot's perspective) extends in the positive y -direction. Furthermore, the center of the rectangular volume enclosing the target is taken to be the origin of the (x, y, z) coordinate system. Finally, we also assume that all of the scattering centers $\{g_n\}$ are contained in this bounding box. With these assumptions, the Nyquist criterion requires that the stepsizes for uniform sampling Δ_X , Δ_Y , and Δ_Z satisfy

$$\Delta_X \leq \frac{1}{L}, \quad \Delta_Y \leq \frac{1}{W}, \quad \text{and} \quad \Delta_Z \leq \frac{1}{H}, \quad (13)$$

where L , W , and H are the length, wingspan, and height, respectively, of the target. In order to translate the constraints on Δ_X , Δ_Y , and Δ_Z into values for Δ_{EL} , Δ_{AZ} , and Δ_f , we take differentials of the relationships in (12) and substitute these results into (13). The presence of f as a scaling factor for several of the differential terms makes it impossible to satisfy the Nyquist criterion for all choices of $f \geq 0$. In order to proceed, we must incorporate f within the angular increments. Also, in the interest of simplifying the system of inequalities, we impose the constraint that the two angular increments be the same. Defining $\Delta \triangleq f \Delta_\theta = f \Delta_\phi$, we can then express (13) as

$$|\sin \theta \cos \phi \Delta_f + \cos(\theta + \phi) \Delta| \leq \frac{c}{2L}, \quad (14)$$

$$|\sin \theta \sin \phi \Delta_f + \sin(\theta + \phi) \Delta| \leq \frac{c}{2W}, \quad (15)$$

$$|\cos \theta \Delta_f - \sin \theta \Delta| \leq \frac{c}{2H}. \quad (16)$$

This set of inequalities is a compact representation of the Nyquist criterion for uniform sampling in frequency-aspect space. To solve for the actual increments, we have to specify the desired trade-off between Δ_f and Δ , since one can be made larger at the expense of the other. As a numeric example, consider the case where $\Delta_f = \Delta$ and $L = \max\{L, W, H\}$. In this case, (14) happens to impose

⁷Building a bistatic RCS database actually amounts to uniformly sampling the function $G(f, \theta^i, \phi^i, \theta^s, \phi^s)$. However, it can be shown that the stepsizes which satisfy the Nyquist criterion in the monostatic scenario are lower bounds for all combinations of incident and scattered directions.

the most stringent requirement on the sampling increment. Numeric evaluation of it yields

$$1.618 \Delta_f \leq \frac{c}{2L}, \quad (17)$$

If we take $c = 3 \times 10^8$ m/s, $f = 100$ MHz, and $L = 15$ m (for a fighter-sized aircraft), we have $\Delta_f \leq 6.18$ MHz. Using the largest increment possible and remembering that we chose $\Delta = \Delta_f$ for this example, the frequency-aspect increments for our RCS database would be $\Delta_f = 6.18$ MHz and $\Delta_{EL} = \Delta_{AZ} = \Delta/f = 3.54^\circ$.

3. KALMAN FILTER-BASED TRACKING

In this section, we present the first of our stochastic filters for FM-band air surveillance. This formulation will be based on the extended Kalman filter (EKF), and it will constitute the benchmark system to which we will compare our particle filter in Section 5. We will define its system model, measurement model, and data association technique, under the assumption that commercial FM broadcasts are the waveform of choice. Since the Kalman prediction and update equations can be found in many textbooks on estimation theory [23, 24], we will simply state the necessary results without derivation.

System Model for the Kalman Filter

If we begin by assuming a linear dynamic system model, the standard discrete-time Kalman state equation takes the form

$$\mathbf{x}_{k+1} = F_k \mathbf{x}_k + \mathbf{u}_k, \quad (18)$$

where \mathbf{x}_k is a vector of kinematic components, \mathbf{u}_k is the process noise, and F_k is our (possibly) time-varying state matrix. Since our multitarget-multisensor application will involve targets which are maneuvering, we adopt a constant-acceleration motion model. Thus, in 3 dimensions, \mathbf{x}_k will contain 9 components,

$$\mathbf{x}_k = [p_{x,k} \ v_{x,k} \ a_{x,k} \ p_{y,k} \ v_{y,k} \ a_{y,k} \ p_{z,k} \ v_{z,k} \ a_{z,k}]', \quad (19)$$

where $p_{x,k}$, $v_{x,k}$, and $a_{x,k}$ denote position, velocity, and acceleration, respectively, along the x -axis at time k . The components for the y and z axes are defined in a similar manner. The Newtonian system matrix then assumes a block diagonal form

$$F = \begin{bmatrix} \tilde{F} & 0 & 0 \\ 0 & \tilde{F} & 0 \\ 0 & 0 & \tilde{F} \end{bmatrix}, \quad (20)$$

where motion along each coordinate axis evolves independently according to

$$\tilde{F} = \begin{bmatrix} 1 & T & T^2/2 \\ 0 & 1 & T \\ 0 & 0 & 1 \end{bmatrix}. \quad (21)$$

In this case, as long as the intersample interval T is a constant, F and \tilde{F} do not vary with k .

For this constant-acceleration dynamic system, the process noise \mathbf{u}_k models changes in the state due to changes in the underlying acceleration increment sequence. We assume that $\{\mathbf{u}_k\}$ is a zero-mean white Gaussian sequence with known covariance

$$E[\mathbf{u}_i \mathbf{u}_j'] = Q \delta(i - j), \quad (22)$$

where δ is the Kronecker delta function and Q is block diagonal

$$Q = \begin{bmatrix} \tilde{Q} & 0 & 0 \\ 0 & \tilde{Q} & 0 \\ 0 & 0 & \tilde{Q} \end{bmatrix}. \quad (23)$$

Since this is a formulation in discrete time, each random acceleration increment acts upon the state for the sample period T . Thus, we can find \tilde{Q} by taking the outer product of $[T^2/2 \ T \ 1]'$. This yields

$$\tilde{Q} = \begin{bmatrix} T^4/4 & T^3/2 & T^2/2 \\ T^3/2 & T^2 & T \\ T^2/2 & T & 1 \end{bmatrix} \sigma_u^2, \quad (24)$$

where σ_u^2 is the variance of the noise sequence modeling the acceleration increment process.

Measurement Model for the Kalman Filter

We begin by defining the notation which will be used for our multisensor observation vector. Define $\mathbf{Z}_{m,k}$ as the set of measurements provided by the m^{th} sensor at time k . Then, the multisensor measurement vector at time k can be defined as the concatenation of all the current scans, or $\mathbf{Z}_k = \{\mathbf{Z}_{1,k}, \dots, \mathbf{Z}_{M,k}\}$, where M is the number of sensors (i.e., FM transmitters) used by the radar system. The measurement model for the standard Kalman filter can then be expressed as

$$\mathbf{Z}_k = H_k \mathbf{x}_k + \mathbf{w}_k, \quad (25)$$

where H_k is allowed to vary with time and $\{\mathbf{w}_k\}$ is the measurement noise sequence. For the specific application we are considering in this paper, FM-band air surveillance, the observation vector \mathbf{Z}_k for the Kalman filter will consist of pairs of delay and Doppler shift measurements. If we begin by considering a single target in a clutter-free environment, we have

$$\mathbf{Z}_k = \{(\tau_{1,k}, d_{1,k}) (\tau_{2,k}, d_{2,k}) \dots (\tau_{M,k}, d_{M,k})\}, \quad (26)$$

where $(\tau_{m,k}, d_{m,k})$ is the delay-Doppler measurement from the m^{th} transmitter at time k .⁸ If we define \mathbf{p}_k and \mathbf{v}_k as the position and velocity vectors, respectively, of the target at time k , then the components of \mathbf{Z}_k are related to these state components by

$$\tau_{m,k} = \frac{\|\mathbf{p}_k - \mathbf{p}_{t_m}\| + \|\mathbf{p}_k - \mathbf{p}_r\|}{c}, \quad (27)$$

$$d_{m,k} \approx \frac{\mathbf{v}_k \cdot \hat{\mathbf{n}}_m(\mathbf{p}_k)}{\lambda_m}, \quad (28)$$

⁸(26) also assumes the probability of detection is unity for all M transmitters.

where $m = 1, \dots, M$, \mathbf{p}_{t_m} is the location of the m^{th} transmitter, λ_m is its wavelength, \mathbf{p}_r is the location of the receiver, and $\|\cdot\|$ is the Euclidean norm. In (28), $\hat{\mathbf{n}}_m(\mathbf{p}_k)$ is the inward normal at \mathbf{p}_k of the ellipsoid with foci $(\mathbf{p}_{t_m}, \mathbf{p}_r)$ and major axis length $c \tau_{m,k}$. This normal can be evaluated as

$$\hat{\mathbf{n}}_m(\mathbf{p}_k) = \frac{\mathbf{p}_{t_m} - \mathbf{p}_k}{\|\mathbf{p}_{t_m} - \mathbf{p}_k\|} + \frac{\mathbf{p}_r - \mathbf{p}_k}{\|\mathbf{p}_r - \mathbf{p}_k\|}. \quad (29)$$

From (27)–(29), we see that the relationship between \mathbf{x}_k and \mathbf{Z}_k for our application is a nonlinear one (i.e., $\mathbf{Z}_k = h_k(\mathbf{x}_k)$). Thus, we will have to use an extended Kalman filter, where the nonlinear measurement function h_k is linearized about the one-step state prediction $\hat{\mathbf{x}}_{k|k-1}$. (The one-step prediction is defined as $\hat{\mathbf{x}}_{k|k-1} \approx E[\mathbf{x}_k | \mathcal{Z}_{k-1}]$, where \mathcal{Z}_{k-1} is the set of all measurements received through time $k-1$, $\mathcal{Z}_{k-1} = \{\mathbf{Z}_1, \dots, \mathbf{Z}_{k-1}\}$, and the approximate equality is due to the nonlinearity of the measurement function.) The resulting Jacobian matrix can be written as

$$\tilde{H}_k \triangleq [h_k(\mathbf{x}) \nabla_{\mathbf{x}}']_{\mathbf{x}=\hat{\mathbf{x}}_{k|k-1}}, \quad (30)$$

where $\nabla_{\mathbf{x}}$ is the gradient operator expressed as a column vector.⁹ Using the definitions of delay and Doppler shift from (27)–(29), we see that \tilde{H}_k will have the following form

$$\tilde{H}_k = \begin{bmatrix} \vdots & & & & & & & & \\ \frac{\partial \tau_m}{\partial p_x} & 0 & 0 & \frac{\partial \tau_m}{\partial p_y} & 0 & 0 & \frac{\partial \tau_m}{\partial p_z} & 0 & 0 \\ \frac{\partial d_m}{\partial p_x} & \frac{\partial d_m}{\partial v_x} & 0 & \frac{\partial d_m}{\partial p_y} & \frac{\partial d_m}{\partial v_y} & 0 & \frac{\partial d_m}{\partial p_z} & \frac{\partial d_m}{\partial v_z} & 0 \\ \vdots & & & & & & & & \end{bmatrix}, \quad (31)$$

where the two rows shown correspond to the delay-Doppler pair from the m^{th} scan, and all partial derivatives are evaluated about $\hat{\mathbf{x}}_{k|k-1}$ as indicated in (30). After some algebra, the following equations for the partial derivatives can be obtained

$$\frac{\partial \tau_m}{\partial \mathbf{p}} = -\frac{\hat{\mathbf{n}}_m(\mathbf{p}_k)}{c}, \quad (32)$$

$$\frac{\partial d_m}{\partial \mathbf{v}} = \frac{\hat{\mathbf{v}}_m(\mathbf{p}_k)}{\lambda_m}, \quad (33)$$

$$\frac{\partial d_m}{\partial \mathbf{p}} = \frac{1}{\lambda_m} \left(\frac{(\mathbf{p}_k - \mathbf{p}_{t_m})(\mathbf{p}_k - \mathbf{p}_{t_m})'}{\|\mathbf{p}_k - \mathbf{p}_{t_m}\|^3} + \frac{(\mathbf{p}_k - \mathbf{p}_r)(\mathbf{p}_k - \mathbf{p}_r)'}{\|\mathbf{p}_k - \mathbf{p}_r\|^3} - D_{m,k} I \right) \cdot \mathbf{v}_k, \quad (34)$$

where the normal vector from (29) has been used, I is the identity matrix, and $D_{m,k}$ is defined as

$$D_{m,k} = \frac{\|\mathbf{p}_k - \mathbf{p}_{t_m}\| + \|\mathbf{p}_k - \mathbf{p}_r\|}{\|\mathbf{p}_k - \mathbf{p}_{t_m}\| \|\mathbf{p}_k - \mathbf{p}_r\|}. \quad (35)$$

Please note that in (32)–(35), the position and velocity vectors would be more accurately written as $\hat{\mathbf{p}}_{k|k-1}$ and

⁹Strictly speaking, we could drop the subscript k from h_k since the locations of the transmitters and receiver do not vary with time for our application.

$\hat{\mathbf{v}}_{k|k-1}$ since they are obtained from the components of the one-step prediction $\hat{\mathbf{x}}_{k|k-1}$. This completes our specification of \tilde{H}_k . A description of the measurement noise sequence $\{\mathbf{w}_k\}$ is much simpler. We assume that it is an independent zero-mean Gaussian process with known covariance

$$E[\mathbf{w}_i \mathbf{w}_j'] = R\delta(i - j), \quad (36)$$

where R is block diagonal

$$R = \begin{bmatrix} \tilde{R}_1 & 0 & \cdots & 0 \\ 0 & \tilde{R}_2 & & \\ \vdots & & \ddots & \\ 0 & & & \tilde{R}_M \end{bmatrix}, \quad (37)$$

and the measurement covariance corresponding to the m^{th} scan is defined as

$$\tilde{R}_m = \begin{bmatrix} \sigma_{\tau_m}^2 & 0 \\ 0 & \sigma_{d_m}^2 \end{bmatrix}. \quad (38)$$

In our experiments, we will assume that $\sigma_{\tau_m}^2$ and $\sigma_{d_m}^2$ do not vary with m .

With these definitions for our system and measurement models, we can now express the prediction and update equations for our extended Kalman filter. Given the (approximate) mean $\hat{\mathbf{x}}_{k-1|k-1}$ and covariance $\Sigma_{k-1|k-1}$ of the state posterior after $k-1$ measurements, the mean and covariance of the one-step state prediction can be computed as

$$\hat{\mathbf{x}}_{k|k-1} = F \hat{\mathbf{x}}_{k-1|k-1}, \quad (39)$$

$$\Sigma_{k|k-1} = F \Sigma_{k-1|k-1} F' + Q, \quad (40)$$

for $k = 1, 2, \dots$, where F is defined by (20)–(21), Q is defined by (23)–(24), $\hat{\mathbf{x}}_{0|0}$ is the mean of the initial state distribution, and $\Sigma_{0|0}$ is its covariance. The measurement update then proceeds according to

$$\hat{\mathbf{x}}_{k|k} = \hat{\mathbf{x}}_{k|k-1} + K_k (\mathbf{Z}_k - h_k(\hat{\mathbf{x}}_{k|k-1})), \quad (41)$$

$$\Sigma_{k|k} = \Sigma_{k|k-1} - K_k S_k K_k', \quad (42)$$

where h_k is defined implicitly by (27)–(29), K_k is the time-varying Kalman gain matrix

$$K_k = \Sigma_{k|k-1} \tilde{H}_k' S_k^{-1}, \quad (43)$$

and S_k is the covariance of the innovation

$$\nu_k = \mathbf{Z}_k - h_k(\hat{\mathbf{x}}_{k|k-1}). \quad (44)$$

The matrix S_k can be calculated using

$$S_k = \tilde{H}_k \Sigma_{k|k-1} \tilde{H}_k' + R, \quad (45)$$

where \tilde{H}_k is defined by (31)–(35) and R is given in (37)–(38).

Data Association for the Kalman Filter

So far, in the development of our extended Kalman filter, we have assumed a tracking task involving a single target in a clutter-free environment with zero probability of miss. A much more realistic scenario would permit multiple targets, clutter returns (i.e., false alarms), and a nonzero probability of miss. Then, our measurement vector \mathbf{Z}_k could very well contain more or less than M delay-Doppler pairs. *Data association* is the process of dealing with the uncertain origin of each measurement in this case. In the presence of clutter, data association involves deciding which measurements were produced by real targets and which are spurious. In the multitarget scenario, the measurements believed to be from actual targets must then be divided among the existing tracks.¹⁰ A *feasible* association of data to targets (and clutter) must satisfy two criteria: 1) no measurement can have more than one target as its source, and 2) in any given sample period, no target can produce more than one set of measurements (i.e., a single delay-Doppler pair per transmitter). A measurement which is not associated with any target is designated as clutter; a target which goes undetected at time k will not be associated with any measurements from \mathbf{Z}_k .

Methods of data association can be grouped into two broad categories: hard associations techniques and soft association techniques [25]. *Hard* association techniques attempt to determine the single “best” association of measurements to targets for use in the Kalman update stage. This optimal association is determined according to some suitably chosen criterion (e.g., maximum likelihood). *Soft* association techniques use all measurements to update all tracks, but the influence of each measurement varies with the probability that it was produced by the target under consideration. We note that both hard and soft approaches to data association ideally involve enumerating all feasible associations (to either select the best one or to estimate their probabilities of occurrence). This can become a daunting prospect in the multitarget-multisensor scenario because the number of associations grows exponentially with the number of targets and the number of sensors. Details about employing a Lagrangian relaxation technique to handle this combinatorial explosion can be found in [26, 27]. In this paper, our simulations involve a 2-target 3-sensor scenario in which it is still possible to enumerate all feasible associations for light to moderate clutter densities. As such, we will use the popular Joint Probabilistic Data Association (JPDA) technique [28] in our extended Kalman filter. Since JPDA is a well-known (soft) association technique, we will only highlight its key equations, drawing attention to details which are specific to our application.

The JPDA algorithm has four basic steps: formation of the validation matrix, enumeration of all feasible associations, calculation of association probabilities, and update of the

¹⁰Measurements believed to be produced by real targets can also be used to initiate new tracks, but we will ignore that possibility in this paper.




Figure 3. Illustration of measurement gates and the corresponding validation matrix.

Kalman state density. The first step, formation of the validation matrix, reduces computation in the subsequent steps. In Fig. 3, a simple example involving a single sensor and two targets is shown. In a slight abuse of earlier notation, the subscripts in the figure index the measurements and targets, not time k or the scan m . We also use a lowercase \mathbf{z} to denote single delay-Doppler pairs. Thus, $\hat{\mathbf{z}}_1$ and $\hat{\mathbf{z}}_2$ are the predicted measurements for targets t_1 and t_2 . $\{\mathbf{z}_1, \dots, \mathbf{z}_4\}$ is the set of delay-Doppler measurements received. The extent of a target's validation gate is chosen so that any measurement falling outside of it is very unlikely to have been produced by the target. This gating eliminates many feasible associations which would have negligibly small probabilities. The gates in Fig. 3 are represented by the two ellipses. The resulting validation matrix, Ω , is shown at the right. Each row of the matrix corresponds to a different measurement while each column corresponds to a different target. An extra column, t_0 , is included to represent the possibility that a measurement is a false alarm. Thus, $\Omega(i, t) = 1$ means that the i^{th} measurement may have originated from the t^{th} target. The need for a joint approach to data association is highlighted whenever a measurement lies within multiple gates. In Fig. 3, we see that \mathbf{z}_2 lies within both gates. Data assignment for the two targets must proceed jointly since association of \mathbf{z}_2 with t_1 precludes its association with t_2 , and vice versa. For our multisensor application, we generate a validation matrix Ω_m , $m = 1, \dots, M$, for the measurements from each transmitter. In the presence of clutter, the measurement counts will differ from sensor to sensor. Furthermore, a target which has been detected in the current scan from one transmitter may be missed in the current scan from another. Thus, while each Ω_m will have the same number of columns, the number of rows will vary with m .

The second step in the JPDA algorithm requires the enumeration of all feasible associations. The j^{th} feasible association of data to targets and clutter, $\chi^{(j)}$, is the union of feasible associations for each transmitter,

$$\chi^{(j)} = \bigcup_{m=1}^M \chi_m^{(j)}. \quad (46)$$

A feasible association for the m^{th} transmitter can be found by changing ones to zeros within Ω_m until each row contains exactly one nonzero entry and each column contains

at most one nonzero entry (except for column t_0 which may contain any number of ones). The computational burden of data association in the multisensor scenario can become overwhelming because any change in one of the sub-assignments $\chi_m^{(j)}$ constitutes a new association $\chi^{(l)}$.

The third step in the JPDA algorithm involves computation of $P(\chi^{(j)}|\mathcal{Z}_k)$ for each feasible association. We simply extend the typical JPDA assumptions [28] to the multisensor case. That is, we assume that each target is detected independently in the scan from the m^{th} transmitter with probability P_{D_m} . If detected by all M sensors, the resulting $2M$ -dimensional innovation is assumed to have a Gaussian distribution with mean $h_k(\hat{\mathbf{x}}_{t,k|k-1})$ and covariance $S_{t,k}$, where the subscripts t identify the target to which this state corresponds.¹¹ The number of clutter returns per transmitter is assumed to be Poisson-distributed and independent from transmitter to transmitter and scan to scan. The clutter locations are assumed to be uniformly distributed in some region of delay-Doppler space which contains the predicted measurements of each target. Finally, in order to avoid complications which arise in the JPDA state update equations when a target is detected by some transmitters but missed by others, we will make one non-standard assumption: for the rest of this paper, we will assume that $P_{D_m} = 1 \forall m$. This means that all innovations will have dimension $2M$, and we will not need to include any correction terms in the standard JPDA covariance update equation. With this last assumption, we are left with a very simple form for the j^{th} association probability,

$$P(\chi^{(j)}|\mathcal{Z}_k) = \frac{1}{C} \prod_{t=1}^{N_t} \frac{\exp\left[-\frac{1}{2} \nu_{t,k}^{(j)'} S_{t,k}^{-1} \nu_{t,k}^{(j)}\right]}{|S_{t,k}|^{1/2}}, \quad (47)$$

where C is a normalization constant, N_t is the number of targets being tracked, and the innovation $\nu_{t,k}^{(j)}$ is formed by taking those components of \mathbf{Z}_k which $\chi^{(j)}$ associates with target t . Note that without our assumption of unity probability of detection, (47) would need to be multiplied by various powers of P_{D_m} and $(1 - P_{D_m})$. We would also need to multiply our expression by a term involving the number of measurements attributed to clutter by $\chi^{(j)}$. Under our assumption of zero probability of miss, these terms become constants and can be absorbed into C .

The fourth and final step of the JPDA algorithm involves the modification of the state update equations, (41)–(42). Using the association probabilities calculated in the previous step, we first compute the weighted innovation

$$\tilde{\nu}_{t,k} = \sum_j P(\chi^{(j)}|\mathcal{Z}_k) \nu_{t,k}^{(j)}, \quad (48)$$

where $t = 1, \dots, N_t$ is the target index, and each sum is taken over all feasible associations. The JPDA update

¹¹ In the presence of multiple measurements, the innovation density will actually be a mixture of Gaussians. However, the nonlinearity of our measurement model makes this an approximation as well.

of the mean of each target's state density can then be expressed as

$$\hat{\mathbf{x}}_{t,k|k} = \hat{\mathbf{x}}_{t,k|k-1} + K_{t,k} \tilde{\nu}_{t,k}, \quad (49)$$

where the rather cumbersome t subscripting is meant to indicate that while the association is conducted jointly across all targets, each target's state density is updated separately. In the case of the covariance update, our expression takes on a slightly simpler form when compared to the standard JPDA algorithm because of our assumption of zero probability of miss,

$$\Sigma_{t,k|k} = \Sigma_{t,k|k-1} - K_{t,k} \cdot (S_{t,k} - \tilde{S}_{t,k}) K_{t,k}' \quad (50)$$

where

$$\tilde{S}_{t,k} = \left(\sum_j P(\chi^{(j)} | \mathcal{Z}_k) \nu_{t,k}^{(j)} \nu_{t,k}^{(j)'} \right) - \tilde{\nu}_{t,k} \tilde{\nu}_{t,k}' \quad (51)$$

is the weighted covariance of the innovation for target t which alters the state covariance to account for the uncertainty in data association.

4. PARTICLE FILTER-BASED TRACKING AND CLASSIFICATION

In the previous section, we presented an extended Kalman filter for an FM-band air surveillance radar. In this section, we will introduce an alternative to the EKF formulation which permits radar cross section to be included in the measurement vector. Before proceeding with definitions for its system model, measurement model, and data association technique, we first provide some motivation for sampling-based approaches to density estimation. This motivation will lead us to the weighted bootstrap and its use for recursive Bayesian inference.

Recursive Bayesian Inference via Weighted Bootstrap

Sampling-based approaches to density estimation are discussed extensively in the statistics literature [1, 2, 29, 30]. We can motivate these approaches by considering the posterior density. In the non-recursive case, Bayes' rule yields

$$p(\mathbf{x} | \mathbf{z}) = \frac{p(\mathbf{z} | \mathbf{x}) p(\mathbf{x})}{\int_{\mathbf{x}} p(\mathbf{z} | \mathbf{x}) p(\mathbf{x}) d\mathbf{x}}. \quad (52)$$

Thus, evaluation of the posterior requires knowledge of both the prior $p(\mathbf{x})$ and the likelihood function $p(\mathbf{z} | \mathbf{x})$, in addition to an integration to find the normalizing constant. Assuming we were able to find $p(\mathbf{x} | \mathbf{z})$, its use in inference would involve integrals of the form

$$E[g(\mathbf{x}) | \mathbf{z}] = \int_{\mathbf{x}} g(\mathbf{x}) p(\mathbf{x} | \mathbf{z}) d\mathbf{x}, \quad (53)$$

for a suitable choice of $g(\cdot)$. Unfortunately, in many cases, the functional form of the prior required in (52) is unknown. Furthermore, even if the prior is available, explicit

evaluation of the integrals in (52) and (53) is rarely possible, requiring sophisticated approximation techniques instead.

Random-sampling approaches are motivated by the essential duality between a sample and the density from which it is drawn [1]. In the limit of an infinite number of independent samples, we could deduce the corresponding density exactly. Furthermore, the Weak Law of Large Numbers enables us to find posterior moments via simple averaging

$$\frac{1}{N} \sum_{i=1}^N g(\mathbf{y}^{(i)}) \xrightarrow{p} E[g(\mathbf{x}) | \mathbf{z}] \quad \text{as } N \rightarrow \infty, \quad (54)$$

where $\{\mathbf{y}^{(1)}, \mathbf{y}^{(2)}, \dots, \mathbf{y}^{(N)}\}$ are independent, identically distributed (iid) samples drawn from $p(\mathbf{x} | \mathbf{z})$ and $E[g^2(\mathbf{x}) | \mathbf{z}]$ is assumed to be finite. Thus, given a set of N iid samples drawn from some prior $p(\mathbf{x})$, upon receipt of a measurement \mathbf{z} , we wish to form a new set of N samples distributed approximately according to the posterior $p(\mathbf{x} | \mathbf{z})$. The weighted bootstrap technique (also known in the literature as factored sampling or importance sampling/resampling) provides us with the solution [1, 5]. The following proposition gives the details in Kalman filter state-space notation.

Proposition 1 *Recursive Bayesian Filtering via Weighted Bootstrap*

Given a set of N samples $\{\mathbf{x}_{k-1|k-1}^{(1)}, \dots, \mathbf{x}_{k-1|k-1}^{(N)}\}$ drawn independently from $p(\mathbf{x}_{k-1} | \mathcal{Z}_{k-1})$, a new set $\{\mathbf{x}_{k|k}^{(1)}, \dots, \mathbf{x}_{k|k}^{(N)}\}$ distributed approximately as $p(\mathbf{x}_k | \mathcal{Z}_k)$ may be obtained by the following two steps.

1) Prediction: Pass each sample $\mathbf{x}_{k-1|k-1}^{(i)}$ through the (possibly nonlinear) system model f_k to yield its one-step prediction

$$\mathbf{x}_{k|k-1}^{(i)} = f_{k-1} \left(\mathbf{x}_{k-1|k-1}^{(i)}, \mathbf{u}_{k-1}^{(i)} \right), \quad (55)$$

where each $\mathbf{u}_{k-1}^{(i)}$ is drawn independently from $p(\mathbf{u}_{k-1})$.

2) Update: Given observation \mathbf{Z}_k , compute a weight for each of the one-step predictions

$$b_i = \frac{p(\mathbf{Z}_k | \mathbf{x}_{k|k-1}^{(i)})}{\sum_{j=1}^N p(\mathbf{Z}_k | \mathbf{x}_{k|k-1}^{(j)})}. \quad (56)$$

Then, draw $\{\mathbf{x}_{k|k}^{(1)}, \dots, \mathbf{x}_{k|k}^{(N)}\}$ from the one-step predictions such that

$$\Pr(\mathbf{x}_{k|k}^{(n)} = \mathbf{x}_{k|k-1}^{(i)}) = b_i \quad (57)$$

for $i, n = 1, 2, \dots, N$.

Thus, we see how the weighted bootstrap technique can




Figure 4. Two coordinate systems: (a) body-centered (b) inertial (receiver-centered).

be applied to the problem of recursive Bayesian filtering.¹² This approach requires three capabilities: 1) the ability to sample from the initial prior $p(\mathbf{x}_0)$, 2) knowledge of the functional form of the likelihood, and 3) the ability to sample from the distribution governing the process noise $p(\mathbf{u}_k)$. It is important to note, in contrast with the standard Kalman filter, this sample-based approach does *not* impose any constraints on the system model or the noise sequences (beyond the three capabilities just enumerated). The system model can be nonlinear, the measurement noise can be non-Gaussian, and most importantly for our purposes, an explicit (closed-form) expression for the measurement model is *not* required.

System Model for the Particle Filter

In this section, we specify the system model which we propose for our particle filter. Since there are no constraints on linearity, we are free to use the correct rigid-body motion model for air flight [31]. The motion of a rigid body has six dynamic degrees of freedom: three to define the location of the object in three-dimensional space and three to define its angular orientation. To express these six quantities, we begin by defining two sets of coordinate axes: body-centered and inertial. Body-centered coordinates are convenient for representing aspects of a target's rotational motion; inertial coordinates are used to express target trajectory. An example of a body-centered coordinate system is depicted in Fig. 4(a). For inertial coordinates, we assume the origin is collocated with our radar receiver with $\hat{i}_{x,I}$ pointing east, $\hat{i}_{y,I}$ pointing north, and $\hat{i}_{z,I}$ pointing up. An example of an inertial coordinate system is provided in Fig. 4(b). Since the orientation of the body-centered axes varies with respect to the inertial frame as the target's attitude changes, let O_k denote the time-varying orthogonal transformation which brings the body-centered axes into alignment with the inertial axes. A vector in the body-centered system is transformed into inertial coordinates by multiplication with O_k .

¹²Strictly speaking, resampling in the update step of Proposition 1 is not necessary. Instead, a cumulative weight can be maintained for each sample, $\tilde{b}_i(k) = b_i \tilde{b}_i(k-1)$. Inference at time index k is then performed using the normalized $\{\tilde{b}_i(k)\}_{i=1}^N$ as sample weights instead of $\frac{1}{N}$. See [30] for further discussion.

It is also convenient to formally define the skew symmetric matrix of body-centered angular velocities

$$\Omega(\boldsymbol{\omega}_k) \triangleq \begin{bmatrix} 0 & -\omega_{z,k} & \omega_{y,k} \\ \omega_{z,k} & 0 & -\omega_{x,k} \\ -\omega_{y,k} & \omega_{x,k} & 0 \end{bmatrix}, \quad (58)$$

where $\boldsymbol{\omega}_k = [\omega_{x,k} \ \omega_{y,k} \ \omega_{z,k}]'$ is the projection of the angular velocity vector on the body-centered coordinate axes (i.e., $\omega_{x,k}$ is the roll rate, $\omega_{y,k}$ is the pitch rate, and $\omega_{z,k}$ is the yaw rate). With O_k and $\Omega(\boldsymbol{\omega}_k)$ defined, we can now present the discretized version of our rigid-body target dynamics.

$$\mathbf{p}_{k+1} = \mathbf{p}_k + T O_k \mathbf{v}_k, \quad (59)$$

$$O_{k+1} = O_k e^{T \Omega(\boldsymbol{\omega}_k)}, \quad (60)$$

$$\mathbf{v}_{k+1} = \mathbf{v}_k + T(\mathbf{f}_k - \Omega(\boldsymbol{\omega}_k)\mathbf{v}_k), \quad (61)$$

where \mathbf{p}_k is inertial position, \mathbf{v}_k is body-centered linear velocity, and \mathbf{f}_k is the normalized force expressed in body-centered coordinates. Note that these difference equations represent only nine of the twelve degrees of freedom needed to completely describe rigid-body dynamics. The missing vector equation relates applied torque to changes in $\boldsymbol{\omega}_k$. However, this fourth equation requires knowledge of a target's moment of inertia about each body-centered axis, information that may not be available for all aircraft of interest. Furthermore, $\boldsymbol{\omega}_k$ can change abruptly during maneuvers, making it very difficult to reliably estimate. For these reasons, we include $\boldsymbol{\omega}_k$ in the process noise instead of the state. In addition, we will also include as noise the total applied force \mathbf{f}_k , which models the effects of thrust, aerodynamics, and gravity. Thus, our process noise is defined as

$$\mathbf{u}_k = \{\mathbf{f}_k, \boldsymbol{\omega}_k\}, \quad (62)$$

where each component is modeled, for sampling during state prediction, as Gaussian distributed about its minimum mean-squared error (MMSE) estimate from the previous sample period. This lends a certain amount of memory to our noise model, an assumption which is increasingly valid as the sample period T decreases. We can now list our complete state vector,

$$\mathbf{x}_k = \{\mathbf{p}_k, \mathbf{v}_k, O_k, \alpha\}, \quad (63)$$

where the discrete-valued target class label α is an element from $\{1, 2, \dots, L\}$, and L is the number of target classes in our library (i.e., the number of targets for which we have computed RCS tables).¹³ With these definitions for \mathbf{x}_k and \mathbf{u}_k , equations (59)–(61) implicitly define the nonlinear system function f_k for our particle filter.

Measurement Model for the Particle Filter

Having discussed the prediction step of our particle filter, we now specify h_k and \mathbf{w}_k for the measurement update. As

¹³Although O_k is a 3×3 matrix, it is completely specified by three angles, such as roll, pitch, and yaw. Thus, the state of our particle filter contains 10 unique components.

we did for the extended Kalman filter, we first consider the simplest case of a single target in a clutter-free environment with zero probability of miss. We will assume that at each sample instant k , our particle filter is provided with M pairs of delay-Doppler measurements (one pair per transmitter). However, because the particle filter does not require an explicit form for h_k , we are free to include one additional type of measurement in \mathbf{Z}_k : RCS. In Section 2, we discussed at length the complex relationship between target position, orientation, and RCS. We saw that for a given \mathbf{p}_k , O_k and target class, evaluation of the resulting RCS required the solution of two integral equations (one to find the induced surface current, the other to find the scattered far-field). In order to incorporate RCS as a measurement in our particle filter, we first compile RCS tables off-line for all target classes of interest. Then, given the pdf of the measurement noise, evaluation of the RCS component of the state likelihood is largely a matter of table look-up.

To evaluate the likelihood that the RCS components of \mathbf{Z}_k were produced by a target with particle state $\mathbf{x}_{k|k-1}^{(i)}$, we need to define that part of the measurement noise process which affects RCS. If we assume that the passive radar system we are designing operates coherently (by making use of the rather limited bandwidth in the FM radio signal), in theory we could recover the complex scattering coefficient $s_{m,k}$ due to transmitter m at time k .¹⁴ However, because of the integration times used in the front-end filter banks, the phases of the scattering coefficients are typically unusable for classification. Thus, to make our simulations more realistic, we will attempt to classify using $|s_{m,k}|$, $m = 1, \dots, M$, only. Since RCS is defined as $\sigma_{m,k} = |s_{m,k}|^2$, this is equivalent to using RCS.¹⁵ Since RCS is a nonnegative quantity, an assumption of Gaussian measurement noise is not appropriate. Instead, we seek to model the thermal noise in the in-phase and quadrature components of the (coherent) matched filter. The specific model we adopt for the RCS component of the measurement noise is

$$\sigma_{m,k} = \left(s_{m,k}^{(R)} + w_1 \right)^2 + \left(s_{m,k}^{(I)} + w_2 \right)^2, \quad (64)$$

where the superscripts identify the real and imaginary components of $s_{m,k}$, and w_1 and w_2 are independent, identically distributed Gaussian random variables with zero mean and variance $\sigma_{s_m}^2$. We will assume that the variance of the thermal noise is independent of the FM signal being received and simplify notation by dropping the subscript m

¹⁴We will assume that the receiver measures horizontal polarization only. Furthermore, since cross-polarization terms are often more than an order of magnitude smaller than co-polarizations (i.e., $|s_{hv}| \ll |s_{hh}|$), we will focus on s_{hh} . For notational simplicity, we omit the subscript hh .

¹⁵Strictly speaking, the receiver measures the power of the in-phase and quadrature components; equation (2) must be used to recover the magnitude of the scattering coefficient. Of course, (2) requires knowledge of the target's position. Since the typical error in our particle filter's position estimate is almost negligible when compared to the distances between the target and the transmitters/receiver, we make the simplifying assumption that $|s_{m,k}|$ or RCS is measured directly.

from $\sigma_{s_m}^2$. The distribution for $\sigma_{m,k}$ is similar to a Rician distribution, except there is no square root in (64). The resulting likelihood can be shown to be

$$p(\sigma_{m,k} | s_{m,k}) = \frac{1}{2\sigma_s^2} \exp \left[-\frac{\sigma_{m,k} + |s_{m,k}|^2}{2\sigma_s^2} \right] \cdot I_0 \left[\frac{\sqrt{\sigma_{m,k}} |s_{m,k}|}{\sigma_s^2} \right], \quad (65)$$

where I_0 is the zero-order modified Bessel function of the first kind. Please note that in (65), we have the unfortunate situation that the Greek letter sigma represents two different quantities. To clear up any confusion, $\sigma_{m,k}$ is the RCS measurement from transmitter m at time k while σ_s^2 is the variance of the thermal noise affecting the measurements of the scattering coefficients $s_{m,k}$, $m = 1, \dots, M$.

In summary, in the single target, clutter-free scenario, the measurement vector for our particle filter will consist of M triplets

$$\mathbf{Z}_k = \{(\tau_{1,k} \ d_{1,k} \ \sigma_{1,k}), \dots, (\tau_{M,k} \ d_{M,k} \ \sigma_{M,k})\}. \quad (66)$$

The relationship between $\tau_{m,k}$, $d_{m,k}$ and the position and velocity components of a given particle state $\mathbf{x}_{k|k-1}^{(i)}$ were already presented for the extended Kalman filter in (27)–(29). The relationship between $\sigma_{m,k}$ and the position and orientation components of $\mathbf{x}_{k|k-1}^{(i)}$ was discussed in Section 2; namely that given \mathbf{p}_k and O_k from the particle state and the locations of transmitter m and the receiver, the body-centered incident and scattered directions can be determined for the target at time k . These body-centered angles may then be used to access an RCS table compiled for the target class identified by the label α from $\mathbf{x}_{k|k-1}^{(i)}$. Thus, for a given particle state, we can either calculate or look-up the expected value of \mathbf{Z}_k . Since we continue to assume that the components of the measurement noise are independent, each state likelihood required for the weighted bootstrap will be the product of its component likelihoods, where delay and Doppler remain Gaussian distributed about their expected values, and RCS is distributed according to (65).

Data Association for the Particle Filter

So far, in the development of our particle filter, we have restricted our attention to a task involving a single target in a clutter-free environment with zero probability of miss. As done in Section 3 for the extended Kalman filter, we now would like to permit multiple targets and clutter returns in our formulation. For the multitarget scenario, we could redefine the particle state \mathbf{x}_k to be a concatenation of all N_t states. Then, each particle would contain kinematic information about all N_t targets. Unfortunately, the required sample size of the particle filter increases (exponentially in the worst case) with the dimension of the state vector. If we wanted to track $N_t = 20$ targets simultaneously, such a “super-state” vector would have 200 components, and the resulting particle count could be prohibitive both in

terms of storage and computation. On the other hand, we could simply maintain a separate particle filter for each target. This would yield an overall filter whose particle count grew linearly with N_t . Unfortunately, in this case, joint data association becomes impossible because each particle state contains information about a single target only (and the likelihood that a given target produced some part of the current observation \mathbf{Z}_k is dependent on how likely it is that the *rest* of the targets or clutter can account for the remaining components of \mathbf{Z}_k). Thus, we see that data association in the case of a separate filter for each target can become extremely difficult for any sizable N_t . For the simulations in Section 5, we will use separate particle filters for each target and a nearest neighbor association rule. Since the particle state of a single target does not contain enough information to compute the likelihood of the entire data vector \mathbf{Z}_k , the likelihood measure used in the resampling stage is recast as

$$p\left(\mathbf{Z}_k \mid \mathbf{x}_{k|k-1}^{(i)}\right) \equiv p\left(\mathbf{z}_{1,k}^{(i)}, \dots, \mathbf{z}_{M,k}^{(i)} \mid \mathbf{x}_{k|k-1}^{(i)}\right) = \prod_{m=1}^M \max_{\mathbf{z}_{m,k} \in \mathcal{Z}_{m,k}} p\left(\mathbf{z}_{m,k} \mid \mathbf{x}_{k|k-1}^{(i)}\right), \quad (67)$$

where $\mathbf{z}_{m,k}$ is a dummy variable corresponding to a single triplet of delay, Doppler, and RCS from the m^{th} scan. In words, the most likely triplet of delay, Doppler, and RCS is chosen from each transmitter's scan and used to compute this pseudolikelihood. The terms $\mathbf{z}_{m,k}^{(i)}$ in the middle of (67) indicate that the nearest neighbor likelihood actually accounts for only M data triplets, not the entire measurement vector \mathbf{Z}_k . Thus, the nearest neighbor rule is a suboptimal hard association technique. We also note that there is no mechanism for preventing the same data triplet from being assigned to multiple target tracks.

MMSE State Estimate and MAP Classification

In this section, we briefly discuss how a collection of N particles, drawn approximately from the posterior density of the state, can be used to estimate quantities such as target class and position. As discussed earlier, the Bayesian recursion for our particle filter is initialized by drawing N independent samples from the initial prior $p(\mathbf{x}_0)$. Since target label α is one of our state components, $p(\mathbf{x}_0)$ will incorporate the *a priori* probability of class α entering the radar's surveillance volume. In the case of a uniform marginal density over the target classes, the initial set of samples would contain roughly N/L particles from each target class, where L is the number of classes in our target library. However, as the recursion proceeds, the resampling performed during measurement updates can change this initial distribution drastically. More specifically, targets whose RCS values lead to poor matches with the observed data can end up with negligible likelihoods and little chance of being selected during resampling. As a practical matter, we force a minimum number of particles from each target class to be selected at each resampling stage. Oth-

erwise, a single noisy RCS measurement could potentially eliminate all particles corresponding to the correct target class.¹⁶

Immediately prior to prediction, the approximate maximum *a posteriori* (MAP) estimate of target class is simply the class which has the most particles. This corresponds to integrating out the target class marginal and selecting the mode. However, this estimate is more accurately performed *before* resampling, since resampling introduces extra random variation within the particle set. Thus, after calculating the normalized likelihoods $\{b_i\}$ according to (56) and (67), the MAP estimate of target class is

$$\hat{\alpha}_k = \arg \max_{\alpha \in \{1, \dots, L\}} \{\eta_k(\alpha)\}, \quad (68)$$

where the particle counts are defined as

$$\eta_k(\alpha) = N \sum_{i=1}^N b_i \mathbf{1}\left(\mathbf{x}_{k|k-1}^{(i)} \in \alpha\right) \quad (69)$$

and the indicator function $\mathbf{1}(\cdot)$ is used to determine if a particle belongs to target class α . The (approximate) MMSE estimate of the current position of the MAP target class may then be found using

$$\hat{\mathbf{p}}_k = \frac{N}{\eta_k(\hat{\alpha}_k)} \sum_{i=1}^N b_i \mathbf{p}_k^{(i)} \mathbf{1}\left(\mathbf{x}_{k|k-1}^{(i)} \in \hat{\alpha}_k\right), \quad (70)$$

where $\mathbf{p}_k^{(i)}$ is the position vector extracted from the state of the i^{th} particle, $\mathbf{x}_{k|k-1}^{(i)}$. According to the Law of Large Numbers, this is an approximation to the expected target position conditioned on the observations \mathcal{Z}_k and the event that $\hat{\alpha}_k$ is the correct class.

5. EXPERIMENTAL RESULTS

Having presented multitarget-multisensor formulations for both an extended Kalman filter and a particle filter, we are now ready to compare the tracking performance of these two systems through Monte Carlo simulations. Our intended application is a passive radar system which uses commercial FM broadcasts for tracking (and classification, in the case of the particle filter). To uniquely locate a target in 3-dimensional space, delay measurements from three different transmitters are needed.¹⁷ The same applies to the unique determination of a target's 3-dimensional velocity vector from measurements of Doppler shift. Of course, in the presence of noise, we could use more than three transmitters to improve tracking performance, but the benefit must be weighed against the increase in computation and hardware that would be required when fielding such a system. As such, we will use the minimum number of transmitters in our simulations.

¹⁶The problem of class extinction could also be addressed by increasing the size of the particle filter, N , but the effect on run-time might be prohibitive.

¹⁷The loci of points which are equidistant from a single transmitter-receiver pair form an ellipsoid. Three such ellipsoids will intersect in one or two points.




Figure 6. Noise-free measurements from each FM transmitter. The first column is delay; the second column is Doppler shift; the third column is RCS. Target-1 is the blue curve in all plots; target-2 is the red curve.




Figure 5. True flight paths and locations of the three FM transmitters and the receiver.

To facilitate analysis of both tracking and classification results, we will consider the simplest multitarget scenario: two aircraft. In order to demonstrate the benefit offered by the inclusion of RCS in the measurement vector, we will consider a “close-target” maneuver. By close-target, we mean that the two aircraft will fly in close proximity with nearly identical headings for a short period of time (approximately 6 seconds). The specific scenario we will consider is shown in Fig. 5. The three FM transmitters are identified in the plot by their FCC call signs.¹⁸ Since the plot

uses a receiver-centered coordinate system, the receiver is located at the origin. The true trajectories of the two targets are drawn in blue (target-1) and red (target-2). In order to identify the beginning and end of each trajectory, we give a short written description of the engagement. Target-1 begins heading due east (along the positive x -axis); target-2 begins heading due west. About 10 seconds into the simulation, target-2 flies up approximately 75 meters behind and 200 meters east of target-1. Then, after another 10 seconds, target-2 breaks away and heads west while target-1 loops back due north. We note that while both of the trajectories are synthetic, we attempted to design them with realistic turn radii and air speeds. The resulting noise-free delay, Doppler, and RCS measurements are shown in Fig. 6. As intended, the true delay and Doppler measurements for the two trajectories are nearly identical between 12 and 18 seconds. Our expectation is that, for the Kalman filter, the two targets will become almost indistinguishable during this period, allowing their tracks to be incorrectly switched after they separate (i.e., target-1 may end up heading west while target-2 loops back north). On the other hand, it is hoped that the particle filter’s use of RCS will enable it to differentiate between the two targets, even when there is considerable ambiguity in their delay-Doppler measurements.

Before presenting the simulation results, we first mention four assumptions that were made for both filters. First, we set $\sigma_r^2 = 10^{-4} \text{ ms}^2$, $\sigma_d^2 = 1 \text{ Hz}^2$, and $T = 0.5$ seconds for all transmitters and simulations. While these are rea-

¹⁸The actual latitudes and longitudes of these three FM transmitters, all

located in Maryland, have been used.

Table 1. Number of Track Switches in 100 EKF Simulations vs. Standard Deviation of Process Noise

σ_u	5	10	15	20	25
switch count	54	37	32	36	38

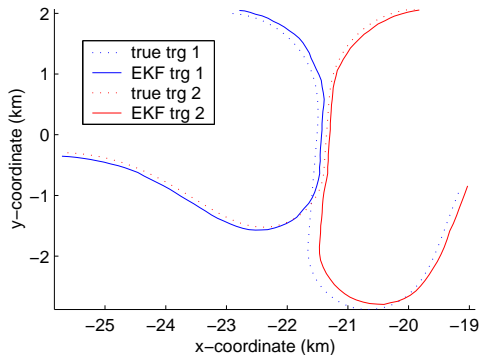


Figure 7. Example of extended Kalman filter switching tracks during a close maneuver.

sonable values for a commercial FM broadcast, they are not intended to correspond to any actual passive radar system. Second, we assumed that both targets maintained a constant altitude of 3000 km above the receiver. Thus, this is a *fixed elevation* scenario. Although both our code and formulation is fully 3-dimensional, the fixed elevation assumption improves performance by reducing the number of state components to six for both filters.¹⁹ Third, we assumed a clutter density of zero throughout, even though the formulation for each filter permits false alarms. This was done primarily for computational reasons, and we expect to have results for tracking/classification in clutter in the near future. Fourth, for the reason mentioned in the data association discussion of Section 3, we set $P_{D_m} = 1$ for all three transmitters.

Kalman Filter Performance

In this section, we present simulation results for the extended Kalman filter. With the four assumptions just listed, the only remaining free parameter for the EKF is the standard deviation of the process noise which models random acceleration increments (σ_u). To determine a good value for this parameter, we ran 100 Monte Carlo simulations for five different choices of σ_u . The results are included in Table 1. The performance metric in this case is the number of track switches (i.e., the number of times the target identities were incorrectly reversed because of their proximity during maneuvers). An example of an EKF track switch is shown in Fig. 7. We find that $\sigma_u = 15$ is the best choice for this task. Smaller values of σ_u increase the chance that the EKF will temporarily lose track of a target at the on-

¹⁹For the particle filter simulations, we assume the target is only allowed to yaw; it is not permitted to pitch or roll. This is not realistic, of course, since an aircraft must roll, to some extent, in order to turn.




Figure 8. Plots of various extended Kalman filter estimates for each target versus their true values (except for the position error plot where the “true” values would be zero). The legend in the lower right pane applies to all four plots.

set of an abrupt maneuver (like target-2 breaking off the engagement midway through its trajectory). Values of σ_u which are too large cause the EKF to rely too heavily on delay-Doppler measurements which, at times, can be very ambiguous. What is surprising in Table 1 is that the best value of σ_u yields a track switch only 32% of the time. As ambiguous as the delay-Doppler measurements appear in Fig. 6, one might expect a track switch percentage closer to 50%. As it turns out, there is still some discriminating power in the Doppler measurements. Although the Doppler plots appear almost identical between 12 and 18 seconds, they actually differ by about 3 Hz. Since $\sigma_d = 1$ Hz, this difference is just above the measurement noise floor. On the other hand, $\sigma_\tau = 0.01$ ms renders the delay measurements almost useless for preventing a track switch.

Next, we consider the performance of the EKF when the integrity of the tracks is maintained. Fig. 8 shows the output from one of the trials that did not switch tracks. The upper right pane is a plot of the resulting position error. What is important to notice here is that the error is relatively constant throughout the maneuver. Because the Doppler measurements are so good, target velocity can be estimated extremely well. Tracks that do not switch tend to maintain the position error bias caused by the random initialization of the filter.²⁰ The bottom two panes in Fig. 8 shown both the true and EKF estimated acceleration components for the targets. We note that other than the sudden onset of target-2’s maneuver around 17 seconds, the EKF does a very good job of tracking the acceleration processes of both targets.

²⁰In each trial, the initial position was assumed to be Gaussian distributed about its true value with a standard deviation of 100 meters along both the x and y axes.




Figure 9. CAD models of the four aircraft used in the simulations. From left to right: VFY-218, Falcon-20, F-22, and Flying Bat.

Particle Filter Performance

In this section, we will explore the effect RCS has on tracking performance. Since RCS is, in part, a function of target geometry, we first must provide identities for our two targets. Fig. 9 depicts CAD models for four different airplanes. The VFY-218 is an experimental aircraft that was never built, but its CAD model is included with the FISC distribution. The Falcon-20 and F-22 are actual airplanes whose CAD models were purchased from Viewpoint, a company which creates models for the entertainment and advertising industries. The fourth model, the Flying Bat, was created here at the University of Illinois for our classification experiments. In all trials, target-1 will be a VFY-218 while target-2 is a Falcon-20. The other two models will be included as confusers within the classifier. Thus, while tracking target-1, our particle filter will also have the task of identifying it from the set {VFY-218, F-22, Flying Bat}. In an analogous fashion, target-2 will have as its class library {Falcon-20, F-22, Flying Bat}. Note that the Falcon-20 cannot be included in target-1’s class library because there would be nothing to stop the particle filter for track 1 from locking onto target-2’s measurements when the two aircrafts’ delay-Doppler measurements became very similar. The same holds for the VFY-218 and target 2’s class library.

Since the computational requirements of the particle filter may substantially exceed those of the EKF, care should be taken in selecting free parameters for the system. The most important parameter to be chosen is N , the number of particles used by the filter. In order to strike a balance between high performance and reasonable computation, our particle filter was implemented with two separate counts: a prediction count and a resample count. The prediction count is the number of particles used in the state prediction stage of the filter. The resample count is the number of particles chosen after likelihoods have been determined during the measurement update. Empirically, we have found for this task that the resample count can be substantially smaller than the prediction count. This is not surprising if we recall that the process noise for our particle filter represents both the total applied force and the angular velocity, $\mathbf{u}_k = \{\mathbf{f}_k, \boldsymbol{\omega}_k\}$. Thus, the prediction count must be high enough to adequately sample the space of all possible forces and angular velocities that a target might suddenly undergo. On the other hand, the resample count only needs to be large

enough to select those particles whose $\{\mathbf{f}_k^{(i)}, \boldsymbol{\omega}_k^{(i)}\}$ pairs resulted in non-negligible likelihoods. Overall, this process of expanding and contracting the particle count is similar to the prior editing procedure discussed in [5]. Before we can discuss the performance for various combinations of prediction and resample counts, we first need to consider the type of error that can occur during joint tracking/classification.

For the specific scenario we considered, there was really only one way that the EKF broke down during simulations: a track switch. The Doppler measurements were of high enough quality that, in a clutter-free environment (with zero probability of miss), it became very unlikely that either track would diverge. The only thing that consistently went wrong was the reversal of the target identities during the ambiguous phase of their trajectories. Because our particle filter must track *and* classify, it has a harder task and is subject to a different type of error. Recall that, in order to increase the robustness of the particle filter, a minimum number of particles must be retained from each class during resampling. This minimum count was set to 2 in all our simulations. For target-1, that meant that at least 2 particles had to be included from the VFY-218, F-22, and Flying Bat classes during resampling. For target-2, the same applied to the Falcon-20, F-22, and Flying Bat. If we consider a prediction count of 1000 and a resampling count of 100, a 2-particle resampling minimum translates into a 20-particle state prediction minimum (since each resampled particle would be “expanded” by a factor of 1000/100). This means that as few as 20 random process noise vectors $\mathbf{u}_k^{(i)}$ will be generated to extend the state estimate of an unlikely target class. Quite simply, this is not enough samples to adequately cover the space of applied forces and angular velocities for a maneuvering aircraft. If an unusually noisy RCS measurement causes the true class to score poorly during the measurement update stage, it may immediately find its share of the particle pool drastically reduced. These required minimum counts are designed to sustain the true class for a few sample periods, but if the correct class does not increase its share of the particle pool during that time, the undersampling of the process noise space will eventually cause its state estimate to diverge. Once the position estimate of the correct class has strayed from the true position of the target, the expected RCS values generated from this erroneous position will no longer match the RCS

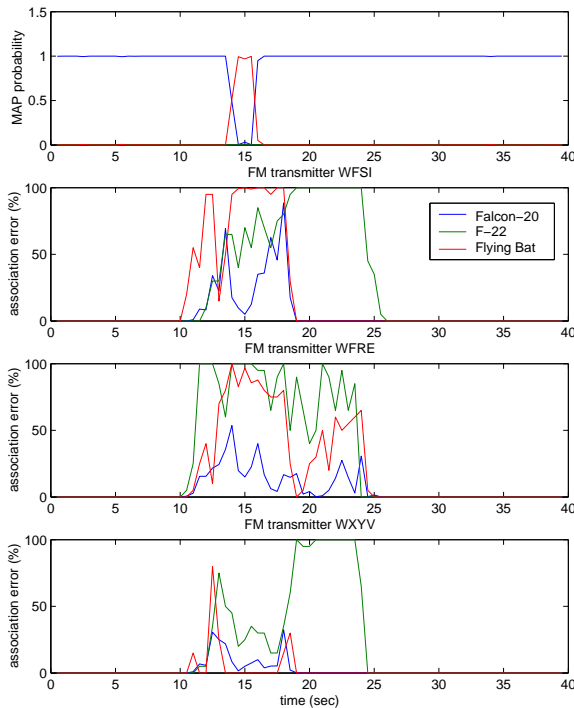


Figure 10. Example where classification jeopardizes successful tracking. The top pane shows the MAP probabilities for all three classes on target-2 data. The lower three panes show the percentage of particles from each class that were incorrectly associated with target-1 data at each update stage.

measurements, and it will be almost impossible for the true class to recover. Thus, it is more likely that an estimated track for the correct class will diverge than switch within the particle filter.

Fig. 10 provides an example from one of the particle filter simulations where the minimum particle counts worked properly. Please note that the colors in this figure represent target class, not track identity. The true class (the Falcon-20 since this is a target-2 example) is depicted in blue while green and red are used for the F-22 and Flying Bat confuser classes, respectively. The top pane shows the evolving MAP probabilities for each class. The lower three panes show, for each FM transmitter, the percentage of particles from each class that were incorrectly associated with target-1 (VFY-218) data by the nearest neighbor rule. We see that even during the ambiguous portion of the trajectory, the majority of Falcon-20 particles (represented by the blue curve) are associated with the correct data from WFRE and WXYV. However, around 15 seconds, the Flying Bat happens to match target-1 RCS measurements from WFSI and WFRE better than the Falcon-20 matches its own (noisy) measurements, and the Flying Bat briefly becomes the MAP class. Without minimum particle counts during resampling, all Falcon-20 particles could have been eliminated at this point. However, a minimum

Table 2. Number of Track Errors and Classification Errors in 50 Simulations vs. Particle Counts

prediction count	resample count	track errors (t_1, t_2)	class errors (t_1, t_2)
1000	100	(4, 0)	(3, 0)
1000	50	(10, 1)	(9, 0)
500	100	(3, 0)	(2, 0)
500	50	(10, 0)	(9, 0)

number is maintained, and this allows the Falcon-20 to regain its proper status as MAP class a few seconds later. This example is meant to highlight two things: 1) the difficulty introduced by simultaneous tracking and classification, and 2) the suboptimality of the nearest neighbor rule. Note that for that brief period when the Flying Bat is the MAP class in Fig. 10, both a target-1 class (the VFY-218) and a target-2 class (the Flying Bat) are being associated with the same measurements. It is precisely this situation which the JPDA algorithm is intended to avoid.

In Table 2, we list the performance of the particle filter using four combinations of particle counts. The performance metrics are somewhat different than the track switch count used in Table 1. The class error column records the number of times in 50 simulations that the MAP class at the end of the run was incorrect for target-1 (t_1) and target-2 (t_2). The track error count records the number of times in these same simulations that the correct class diverged from the true trajectory. As just discussed above, a classification error virtually guarantees that the true class will diverge because of prolonged undersampling of the process noise during state prediction. By comparing the performance of the two different resampling counts, we see that 50 particles do not convey enough information about the posterior for this task. However, an expansion factor of 5 is just as good as 10 during state prediction (compare rows 1 and 3 in Table 2). This is significant because the likelihood computation in the measurement update is one of the most computationally expensive subroutines of the particle filter, and row 3 cuts this computation in half compared to row 1. Overall, if we consider the loss of either track as resulting in an error for the simulation, the best choice of particle counts maintains track 94% of the time. This is to be compared to the best EKF which was successful only 68% of the time. Thus, we have the encouraging result that not only does RCS permit target classification, but it also has the potential to substantially improve the tracking performance by maintaining target identity. This result is even more significant when we consider that our implementation of the particle filter was hindered by a poor data association technique (nearest neighbor), as opposed to the EKF which used the JPDA algorithm.

Fig. 11 shows plots for several components of the particle filter state for one of the Monte Carlo runs. As was the

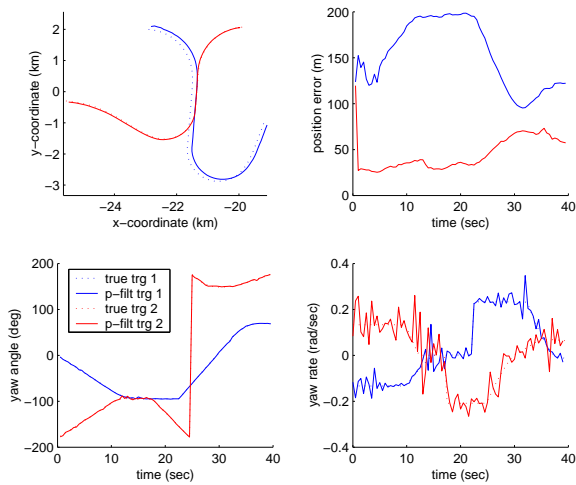


Figure 11. Plots of various particle filter estimates for each target versus their true values (except for the position error plot where the “true” values would be zero). The legend in the lower left pane applies to all four plots.

case with the EKF, the position error plot is largely a function of the random initialization of the particle filter. The delay measurements are so noisy that they have very little effect on reducing the position error bias. Instead, the particle filter concentrates on matching the two higher quality measurements: Doppler shift and RCS. The bottom two panes of Fig. 11 display the filter’s estimates of the yaw angles (orientations) and yaw rates (angular velocities) for both targets. We note that the inclusion of orientation in the EKF tracker would require a nonlinear system model, in addition to the nonlinear measurement model. More importantly, without access to a measurement which varies with angular orientation (such as RCS), an attitude component would be unobservable to the EKF anyway. We see that the RCS data allows the particle filter to maintain a good estimate of both of the targets’ orientations.²¹ Since the angular velocity (i.e., yaw rate in our fixed elevation scenario) is actually a component of the process noise, we expect its estimate to be noise-like in appearance. This is exactly what we see in the lower right pane of Fig. 11.

In Fig. 12, we plot the MMSE estimates of RCS for each of the FM transmitters from one of the simulations. For these plots, we zoom in on the ambiguous region of the trajectories in order to view the fine detail of the estimates. Similar to earlier figures, the dotted lines represent actual values of the parameter being plotted while the solid lines depict the particle filter’s estimate of these quantities. Overall, we see that the particle filter does an admirable job of estimating RCS, a fact which is not surprising when we consider the filter’s ability to track target orientation. How-

²¹ Within our implementation, process noises $\mathbf{u}_k^{(i)}$ were chosen such that there was never more than 5° between the nose of a target and its velocity vector. This is a reasonable limit on the angle of sideslip which amounts to us prohibiting targets from flying sideways or backwards.




Figure 12. Particle filter estimates of RCS for a 20-second segment of the trial. The legend in the second plot applies to all three.

Table 3. Number of Errors in 50 Simulations vs. Standard Deviation of Scattering Coefficient Noise

σ_s	track errors (t_1, t_2)	class errors (t_1, t_2)
0.3	(4, 0)	(3, 0)
0.5	(5, 1)	(5, 0)

ever, Fig. 12 once again demonstrates the inappropriateness of the nearest neighbor association rule for this task. The most striking example of this is in the lower pane (WXYV) at 12 seconds. We see that the target-2 estimate (solid red) has a sharp peak where it suddenly matches the target-1 estimate (solid blue) instead of following the actual target-2 data (dashed red line). This is another example of both targets being allowed to associate with the same measurements. Similar examples can also be found between 12 and 15 seconds in the plots for WFSI and WFRE.

The last thing we briefly mention is our choice of σ_s , the standard deviation of the Gaussian noise process affecting the scattering coefficients. The relationship between this component of the measurement noise and RCS is defined in (64). All simulations presented so far assumed $\sigma_s = 0.3$, which leads to relatively clean RCS measurements. Table 3 contains a comparison between a second choice for σ_s conducted over 50 Monte Carlo simulations. In both cases, the prediction count was 1000, and the resample count was 100. We see little change in performance because of the increase in σ_s . However, it is clear from Fig. 6 that if the noise affecting RCS measurements was on the order of the differences between the RCS curves during the ambiguous

portion of the trajectories, tracking performance could deteriorate rapidly. More work needs to be done to determine both the value of σ_s at which this occurs and a realistic operating range for this parameter based on actual passive radar systems.

6. CONCLUSIONS

We have presented two stochastic filters for an FM-band passive air surveillance radar. The first system used an extended Kalman filter, delay-Doppler measurements, and the JPDA algorithm to track targets. The second system used a particle filter to simultaneously track and classify targets using delay, Doppler, and RCS. At the beginning of the paper, we discussed the complex relationship that exists between target aspect and radar cross section. Because there is no closed-form expression for this relationship, the extended Kalman filter is unable to take advantage of the classification potential of RCS measurements. In addition to our descriptions of both filters, we also presented a significant number of simulation results demonstrating the benefits of RCS and a joint approach to tracking and classification. We considered a challenging 2-target 3-sensor task in which the actual target trajectories nearly coincided for a portion of the total track length. We found that our particle filter was able to maintain track for both targets 94% of the time as compared to 68% using the extended Kalman filter. This result is even more promising when we consider that our particle filter was hindered by a highly suboptimal nearest neighbor data association technique.

In the future, we plan on extending our work to consider scenarios involving clutter and missed detections. We also hope to improve the particle filter's data association capability.

7. ACKNOWLEDGEMENTS

This work was supported in part by a grant from DARPA under Contract AFOSR F49620-98-1-0498. The authors are grateful to Dr. Aaron Lanterman for the many insightful comments he made to improve this paper.

REFERENCES

- [1] A. F. M. Smith and A. E. Gelfand, "Bayesian statistics without tears: A sampling-resampling perspective," *The American Statistician*, vol. 46, pp. 84–88, May 1992.
- [2] A. E. Gelfand and A. F. M. Smith, "Sampling-based approaches to calculating marginal densities," *Journal of the American Statistical Association*, vol. 85, pp. 398–409, June 1990.
- [3] M. I. Miller, A. Srivastava, and U. Grenander, "Conditional-mean estimation via jump-diffusion processes in multiple target tracking/recognition," *IEEE Transactions on Signal Processing*, vol. 43, pp. 2678–2690, November 1995.
- [4] A. D. Lanterman, M. I. Miller, and D. L. Snyder, "General Metropolis-Hastings jump diffusions for automatic target recognition in infrared scenes," *Optical Engineering*, vol. 36, pp. 1123–1137, April 1997.
- [5] N. J. Gordon, D. J. Salmond, and A. F. M. Smith, "Novel approach to nonlinear/non-gaussian Bayesian state estimation," *IEE Proceedings-F*, vol. 140, pp. 107–113, April 1993.
- [6] N. Gordon and A. Whitby, "A Bayesian approach to target tracking in the presence of glint," in *SPIE*, (Orlando, FL), 1995.
- [7] M. Isard and A. Blake, "Condensation — conditional density propagation for visual tracking," *International Journal of Computer Vision*, vol. 29, no. 1, pp. 5–28, 1998.
- [8] D. Avitzour, "Stochastic simulation Bayesian approach to multitarget tracking," *IEE Proceedings-F*, vol. 142, pp. 41–44, April 1995.
- [9] A. Srivastava, "Bayesian filtering for tracking pose and location of rigid targets," in *SPIE Aerosense*, (Orlando, FL), April 2000.
- [10] P. E. Howland, "Target tracking using television-based bistatic radar," *IEE Proceedings-Radar, Sonar and Navigation*, vol. 146, pp. 166–174, June 1999.
- [11] J. Baniak *et al.*, "Silent sentry™ passive surveillance," tech. rep., Lockheed Martin Mission Systems, Gaithersburg, MD, June 1999.
- [12] H. D. Griffiths and N. R. W. Long, "Television-based bistatic radar," *IEE Proceedings-F*, vol. 133, pp. 649–657, December 1986.
- [13] A. D. Lanterman, "Tracking and recognition of airborne targets via commercial television and FM radio signals," in *SPIE Acquisition, Tracking, and Pointing*, (Orlando, FL), April 1999.
- [14] A. W. Rihaczek, *Principles of High-Resolution Radar*. New York, NY: McGraw-Hill, 1969.
- [15] J. L. Eaves and E. K. Reedy, eds., *Principles of Modern Radar*. New York, NY: Van Nostrand Reinhold, 1987.
- [16] J.-S. Chen, "Relative amplitude and phase features for resonance-region radar target identification," in *Proceedings of the IEEE National Aerospace and Electronics Conference*, pp. 97–101, 1989.
- [17] A. A. Ksienski, Y.-T. Lin, and L. J. White, "Low-frequency approach to target identification," *Proceedings of the IEEE*, vol. 63, pp. 1651–1660, December 1975.
- [18] R. F. Harrington, *Field Computation by Moment Methods*. New York, NY: MacMillan, 1968.
- [19] S. M. Rao, D. R. Wilton, and A. W. Glisson, "Electromagnetic scattering by surfaces of arbitrary shape,"

IEEE Transactions on Antennas and Propagation, vol. 30, pp. 409–418, May 1982.

- [20] J. M. Song and W. C. Chew, “Multilevel fast-multipole algorithm for solving combined field integral equations of electromagnetic scattering,” *Microwave and Optical Technology Letters*, vol. 10, pp. 14–19, September 1995.
- [21] Center for Computational Electromagnetics, University of Illinois and Demaco Inc., Champaign-Urbana, IL, *User’s Manual for FISC (Fast Illinois Solver Code)*, October 1996.
- [22] L. C. Trintinalia, R. Bhalla, and H. Ling, “Scattering center parameterization of wide-angle backscattered data using adaptive gaussian representation,” *IEEE Transactions on Antennas and Propagation*, vol. 45, pp. 1664–1668, November 1997.
- [23] A. Gelb, ed., *Applied Optimal Estimation*. Cambridge, MA: M.I.T. Press, 1974.
- [24] H. V. Poor, *An Introduction to Signal Detection and Estimation*. New York, NY: Springer, 2nd ed., 1994.
- [25] Y. Bar-Shalom and X.-R. Li, *Multitarget-Multisensor Tracking: Principles and Techniques*. Storrs, CT: YBS Publishing, 1995.
- [26] K. Pattipati, S. Deb, Y. Bar-Shalom, and J. Robert Washburn, “A new relaxation algorithm and passive sensor data association,” *IEEE Transactions on Automatic Control*, vol. 37, pp. 198–213, February 1992.
- [27] S. Deb, M. Yeddanapudi, K. Pattipati, and Y. Bar-Shalom, “A generalized S - D assignment algorithm for multisensor-multitarget state estimation,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 33, pp. 523–536, April 1997.
- [28] T. E. Fortmann, Y. Bar-Shalom, and M. Scheffe, “Sonar tracking of multiple targets using joint probabilistic data association,” *IEEE Journal of Oceanic Engineering*, vol. OE-8, pp. 173–184, July 1983.
- [29] W. K. Hastings, “Monte Carlo sampling methods using Markov chains and their applications,” *Biometrika*, vol. 57, no. 1, pp. 97–109, 1970.
- [30] J. S. Liu and R. Chen, “Sequential Monte Carlo methods for dynamic systems,” *Journal of the American Statistical Association*, vol. 93, pp. 1032–1044, September 1998.
- [31] B. Friedland, ed., *Control System Design: An Introduction to State-Space Methods*. McGraw-Hill, 1986.



Shawn Herman (S'99) was born in Evanston, IL in 1972. He received the B.S. and M.S. degrees in electrical engineering from the University of Illinois in 1994 and 1996, respectively. From 1996 to 1998, he worked as a Member of Technical Staff in the Speech Recognition Group at Lucent Technologies, Naperville, IL. In 1998, he returned to the University of Illinois to pursue his Ph.D. During the summers of 1999 and 2000, he interned with the Passive Coherent Location Department of Lockheed Martin Mission Systems, Gaithersburg, MD. His research interests include speech recognition, RCS modeling, passive radar systems, and joint tracking/classification.



Pierre Moulin (S'89-M'90) was born in Mons, Belgium in 1963. He received the degree of Ingénieur civil électricien from the Faculté Polytechnique de Mons, Belgium in 1984, and the M.Sc. and D.Sc. degrees in electrical engineering from Washington University in St. Louis in 1986 and 1990, respectively. After working for 5 years as a Research Scientist for Bell Communications Research in Morristown, NJ, he joined the University of Illinois where he is currently an Associate Professor in the Department of Electrical and Computer Engineering. His fields of interest include image and video processing, statistical signal processing and modeling, nonparametric function estimation, and information hiding.