

A Sequential Monte Carlo Framework for Extended Object Tracking

Jaco Vermaak, Norikazu Ikoma and Simon J. Godsill

Abstract

In this paper we consider the problem of extended object tracking. An extended object is modelled as a set of point features in a target reference frame. The dynamics of the extended object are formulated in terms of the translation and rotation of the target reference frame relative to a fixed reference frame. This leads to realistic, yet simple, models for the object motion. We assume that the measurements of the point features are unlabelled, and contaminated with a high level of clutter, leading to measurement association uncertainty. Marginalising over all the association hypotheses may be computationally prohibitive for realistic numbers of point features and clutter measurements. We present an alternative approach within the context of particle filtering, where we augment the state with the unknown association hypothesis, and sample candidate values from an efficiently designed proposal distribution. This proposal elegantly captures the notion of a soft gating function. We demonstrate the performance of the algorithm on a challenging synthetic tracking problem, where the ground truth is known, in order to compare between different algorithms.

Index Terms

Extended object tracking, sequential Monte Carlo methods, particle filters, data association.

I. INTRODUCTION

Traditional methods for radar and sonar based tracking model the target as a point source. This approximation is sufficient for low resolution sensors, or for targets in the far field of the sensor. For high resolution sensors, or objects in the near field of the sensor, the sensor may be able to resolve a number of features on the target. These may provide valuable information about the target motion and orientation. Thus, instead of combining the target object features into a single feature, such as an average or centroid measurement, a richer representation may be achieved by passing all the available features to the tracking algorithm. We will refer to a target observed in this way as an “extended object”.

As in [1] we model an extended object as a set of point features, or highlights, in a target reference frame. Each feature is associated with a detection probability to capture the notion that features do not always generate measurements. We also take account of the fact that the exact locations of the features within the target reference

J. Vermaak and S. J. Godsill are with the Signal Processing Laboratory in the Cambridge University Engineering Department; N. Ikoma is with the Kyushu Institute of Technology in Fukuoka, Japan.

frame are uncertain. We model the dynamics of the extended object in terms of the translation and rotation of the target reference frame relative to a fixed reference frame. In this way we are able to construct realistic, yet computationally tractable, models for the extended object motion.

In most practical tracking applications the sensors yield unlabelled measurements of the point features. Additional clutter measurements may arise due to multi-path effects, the presence of other objects in the sensor range, sensor errors, *etc.* The unlabelled measurements and the presence of clutter lead to a difficult data association problem. Keeping track of all the possible association hypotheses over time, as is the case for the Multiple Hypotheses Tracker (MHT) [2], leads to an NP-hard problem, since the number of association hypotheses grows exponentially over time. Thus, methods are required to reduce the computational complexity. The nearest neighbour filter [3] associates each target with the closest measurement in the target space. However, this simple procedure prunes away many feasible hypotheses. In this respect the Joint Probabilistic Data Association Filter (JPDAF) [3], [4] is more appealing. At each time step improbable hypotheses are pruned away using a gating procedure. A filtering estimate is then computed for each of the remaining hypotheses, and combined in proportion to the corresponding posterior hypothesis probabilities. The main shortcoming of the JPDAF is that, to maintain tractability, the final estimate is collapsed to a single Gaussian, thus discarding much pertinent information. Subsequent work addressed this shortcoming by proposing strategies to instead reduce the number of mixture components in the original mixture to a tractable level [5], [6]. Nevertheless, many feasible hypotheses may be discarded by the pruning mechanisms. The Probabilistic Multiple Hypotheses Tracker (PMHT) [7], [8] (wrongly) assumes the association variables to be independent in order to work around the problems with pruning. It leads to an incomplete data problem that can be efficiently solved using the Expectation Maximisation (EM) algorithm [9]. However, the PMHT is a batch strategy, and thus not suitable for online applications. For further discussion on shortcomings and extensions to the PMHT the reader is referred to [10]. Further developments of data association in the context of MHT are reported in [11]–[13].

The methods discussed above are mostly applicable to linear Gaussian system and observation models. For models with weak non-linearities it is possible to obtain similar algorithms based on approximations such as the Extended Kalman Filter (EKF) [14]. However, the performance of the resulting algorithms degrades rapidly as the non-linearities become more severe. These methods are also not robust to non-Gaussian noise in the system and observation models. These shortcomings have been acknowledged before, and led to the development of numerous strategies to cope with non-linear and non-Gaussian models. One of the first methods to deal particularly with non-Gaussian models is the Gaussian Sum Filter [15] that works by approximating the non-Gaussian target distribution with a mixture of Gaussians. It suffers, however, from the same shortcoming as the EKF in that linear approximations are required. It also leads to a combinatorial growth in the number of mixture components over time, calling for ad-hoc strategies to prune the number of components to a manageable level. An alternative method for non-Gaussian models that does not require any linear approximations has been proposed in [16]. It approximates the non-Gaussian state numerically with a fixed grid, and applies numerical integration for the prediction step and Bayes' theorem for the filtering step. However, the computational cost of the numerical integration grows exponentially with the

dimension of the state-space, and the method becomes impractical for dimensions larger than four.

As an alternative method for general non-linear and/or non-Gaussian models, Particle Filtering [17], [18], also known as Sequential Monte Carlo (SMC) [19]–[21], or CONDENSATION [22], has become a practical numerical technique to approximate the Bayesian tracking recursions. This is due to its efficiency, simplicity, flexibility, ease of implementation, and modelling success over a wide range of challenging applications. It represents the target distribution with a set of samples, or particles, and associated importance weights, which are then propagated through time to give approximations of the target distribution at subsequent time steps. It requires only the definition of a suitable proposal distribution from which new particles can be generated, and the ability to evaluate the system and observation models. As opposed to the strategy in [16], the computational complexity for particle filters is independent of the dimension of the state-space: it grows linearly in the number of particles, with the error decreasing as the square root of the number of particles. For further discussions on the relation between dimensionality and Monte Carlo error the reader is referred to [23], [24].

The data association problem has also been considered within the context of particle filtering. Methods that combine particle techniques with the philosophy behind the JPDAF are described in [25], [26]. In [27] a method is described that computes the distribution of the association hypotheses using a Gibbs sampler [28] at each time step. The method is similar in spirit to the one described in [29] that uses Markov Chain Monte Carlo (MCMC) techniques [30] to compute the correspondences between image points within the context of stereo reconstruction. The main problem with the MCMC strategies is that they are iterative in nature and take an unknown number of iterations to converge. They are thus not very suitable for online applications. In [31] a method is presented where the associations are simulated from an optimally designed importance distribution. The method is intuitively appealing since the association hypotheses are treated in a similar fashion to the target states, so that the resulting algorithm is non-iterative. It is, however, restricted to point targets in the framework of Jump Markov Linear Systems (JMLS) [32].

In this paper we propose an alternative strategy to solve the data association problem within the context of particle filtering. It is similar in spirit to the method described in [31], in that we augment the system state with the unknown association vector. To generate candidate samples for the association vector we construct an efficient proposal distribution based on the notion of a soft gating of the measurements. As opposed to the method in [31] our approach is generally applicable. Its use extends beyond the sequential estimation framework, and it can be applied in any setting, static or dynamic, where the data association problem arises within a probabilistic modelling framework. The framework we present here is a development of our work detailed in [1], [33]. For other recent related work on SMC methods within our research team see *e.g.* [34]–[38].

The remainder of the paper is organised as follows. In Section II we describe the extended object model, and propose a generic model for its motion, based on translation and rotation relative to a fixed reference frame. We outline the measurement process in Section III, and develop a likelihood model for the case where the measurement association is known. We also show how this likelihood can be marginalised over the unknown associations by defining a suitable prior for the associations. In Section IV we formulate the tracking problem within a particle

filtering framework. In addition to the state, we generate samples for the unknown associations by sampling from an efficient proposal distribution, which we derive in Section V. We evaluate the performance of the algorithm on a challenging synthetic tracking problem in Section VI, before concluding with a summary and some remarks in Section VII.

II. TARGET MODEL AND DYNAMICS

As in [1] we model an extended object as a set of fixed point features in a target reference frame T . We will denote these point features by $\mathbf{P}^T = \{\mathbf{p}_i^T\}_{i=1}^{N_p}$, with $\mathbf{p}_i^T = (x_i^T, y_i^T, z_i^T)$ the Cartesian coordinates of the i -th feature point. Note that the extended target is assumed to be rigid, and the number of point features N_p is assumed to be fixed and known. The location and orientation of the target is tracked relative to a fixed reference frame R . For this purpose we define the target state as $\mathbf{x} = (\mathbf{P}^R, \mathbf{t}, \boldsymbol{\theta}, \boldsymbol{\delta})$, where $\mathbf{t} = (x, y, z)$ is the origin of the target frame in the fixed reference frame, $\boldsymbol{\theta} = (\alpha, \beta, \gamma)$ are the roll, pitch and yaw angles of the target frame, measured anti-clockwise around the x , y and z axes, respectively, $\boldsymbol{\delta} = (\boldsymbol{\delta t}, \boldsymbol{\delta \theta})$ are the velocities, with $\boldsymbol{\delta t} = (\delta x, \delta y, \delta z)$ the linear velocity of the origin, and $\boldsymbol{\delta \theta} = (\delta \alpha, \delta \beta, \delta \gamma)$ the angular velocities around the target reference frame axes, and $\mathbf{P}^R = \{\mathbf{p}_i^R\}_{i=1}^{N_p}$ are the positions of the point features in the fixed reference frame R . This full description of the state will not be required in all applications. Depending on the characteristics of the extended object and its motion, some of the state components may become redundant.

With the definition of the state as above any Cartesian point \mathbf{p}^T in the target reference frame can be mapped to a point in the fixed reference frame \mathbf{p}^R through the transformation

$$\begin{aligned} \mathbf{p}^R &= \mathbf{M}(\boldsymbol{\theta})\mathbf{p}^T + \mathbf{t} \\ &= \mathbf{R}(\alpha)\mathbf{P}(\beta)\mathbf{Y}(\gamma)\mathbf{p}^T + \mathbf{t}, \end{aligned} \quad (1)$$

with

$$\mathbf{R}(\alpha) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \alpha & -\sin \alpha \\ 0 & \sin \alpha & \cos \alpha \end{bmatrix} \quad \mathbf{P}(\beta) = \begin{bmatrix} \cos \beta & 0 & \sin \beta \\ 0 & 1 & 0 \\ -\sin \beta & 0 & \cos \beta \end{bmatrix} \quad \mathbf{Y}(\gamma) = \begin{bmatrix} \cos \gamma & -\sin \gamma & 0 \\ \sin \gamma & \cos \gamma & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (2)$$

the rotation matrices corresponding to roll, pitch and yaw, respectively.

We assume the system dynamics to be Markovian and of the form

$$\begin{aligned} p(\mathbf{x}_k | \mathbf{x}_{k-1}) &= p(\mathbf{P}_k^R, \mathbf{t}_k, \boldsymbol{\theta}_k, \boldsymbol{\delta}_k | \mathbf{P}_{k-1}^R, \mathbf{t}_{k-1}, \boldsymbol{\theta}_{k-1}, \boldsymbol{\delta}_{k-1}) \\ &= p(\mathbf{P}_k^R | \mathbf{t}_k, \boldsymbol{\theta}_k) p(\mathbf{t}_k | \mathbf{t}_{k-1}, \boldsymbol{\theta}_k, \boldsymbol{\delta t}_k) p(\boldsymbol{\theta}_k | \boldsymbol{\theta}_{k-1}, \boldsymbol{\delta \theta}_k) p(\boldsymbol{\delta}_k | \boldsymbol{\delta}_{k-1}), \end{aligned} \quad (3)$$

where k denotes the discrete time index. We model the velocities as independent first order Gaussian random walk models, *i.e.*

$$p(\boldsymbol{\delta}_k | \boldsymbol{\delta}_{k-1}) = \mathcal{N}(\boldsymbol{\delta}_k | \boldsymbol{\delta}_{k-1}, \boldsymbol{\Lambda}_\delta) \quad (4)$$

where $\mathcal{N}(\cdot | \boldsymbol{\mu}, \boldsymbol{\Sigma})$ denotes the multi-variate Gaussian distribution with mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$, and $\boldsymbol{\Lambda}_\delta = \text{diag}(\sigma_x^2, \sigma_y^2, \sigma_z^2, \sigma_\alpha^2, \sigma_\beta^2, \sigma_\gamma^2)$ is the diagonal matrix with the variances of the random walk components, which are

assumed to be fixed and known. Other models for the velocities can be straightforwardly accommodated within our framework. Given the velocities the components for the location and orientation of the target frame in (3) can be computed deterministically as

$$\begin{aligned}\boldsymbol{\theta}_k &= \boldsymbol{\theta}_{k-1} + \Delta_T \boldsymbol{\delta} \boldsymbol{\theta}_k \\ \mathbf{t}_k &= \mathbf{t}_{k-1} + \Delta_T \mathbf{M}(\boldsymbol{\theta}_k) \boldsymbol{\delta} \mathbf{t}_k,\end{aligned}\tag{5}$$

where Δ_T is the time step. The first component in (3) is used to express the uncertainty about the exact locations of the point features. We assume it to be of the form

$$p(\mathbf{P}_k^R | \mathbf{t}_k, \boldsymbol{\theta}_k) = \prod_{i=1}^{N_p} \mathcal{N}(\mathbf{p}_{i,k}^R | \widehat{\mathbf{p}}_{i,k}^R, \sigma_{\mathbf{p}}^2 \mathbf{I}_3),\tag{6}$$

where \mathbf{I}_n denotes the $n \times n$ identity matrix, and the mean of the Gaussian distribution for the i -th point feature follows from (1) as $\widehat{\mathbf{p}}_{i,k}^R = \mathbf{M}(\boldsymbol{\theta}_k) \mathbf{p}_i^T + \mathbf{t}_k$. Note that the uncertainty is assumed to be isotropic around the mean location, with a fixed and known variance $\sigma_{\mathbf{p}}^2$.

This completes the specification of the general form of the model. Depending on the characteristics of the object of interest and its motion, the model can often be further restricted, as is exemplified below.

Example 1: In the tracking application we will consider later, we will be interested in tracking a stick like extended object. The stick object coincides with the x axis in the target reference frame, with point features at $\mathbf{p}_1^T = (50, 0, 0)$, $\mathbf{p}_2^T = (-50, 0, 0)$ and $\mathbf{p}_3^T = (20, 0, 0)$. The object is allowed to translate along its x axis, pitch around its y axis, and yaw around its z axis. The velocities for the other components effectively disappear from the state. An example trajectory, simulated from the dynamics in (3), is depicted in Figure 1. For this trajectory the fixed parameters of the model were set to $(\sigma_x, \sigma_\beta, \sigma_\gamma, \sigma_{\mathbf{p}}) = (1, 0.1^\circ, 0.2^\circ, 0)$. ■




Fig. 1. **Example 3D trajectory for the stick object.** Even though no explicit attempt is made to model any physical attributes of the object, the dynamic model generates realistic and feasible trajectories.

III. MEASUREMENTS AND LIKELIHOOD

The likelihood model we describe in this section follows closely the multi-hypothesis likelihood derived in [39]. The extended object tracking is performed from some observer location in the fixed reference frame R . We will denote the observer location by $\mathbf{p}^O = (x^O, y^O, z^O)$, and allow it to vary with time. At each time step¹ the sensors yield M observations $\mathbf{y} = (\mathbf{y}_1 \cdots \mathbf{y}_M)$, where M may also vary with time. The nature of the individual measurements will depend on the characteristics of the sensors. Typically each measurement will correspond to an estimated line of sight from the observer to a point feature. Measurements arise not only from the point features of the extended object. Additional clutter measurements may result due to multi-path effects, other objects within the sensor range, sensor errors, *etc.* In the absence of specific models for clutter and multi-path, we will assume that all of these effects can be modelled within a random clutter framework. In cases where models are available for multi-path, these could be incorporated into our non-linear filtering framework. Hence we will assume that point features on the extended object can generate at most one measurement each at a particular time step, but may also go undetected. We will further assume that several or all of the measurements may be due to unstructured, random clutter.

For a given vector of M measurements we introduce the association hypothesis $\lambda = (\mathbf{r}, M_C, M_T)$, where M_C is the number of clutter measurements, M_T is the number of measurements resulting from the extended object point features, with $M = M_C + M_T$. The elements of the association vector $\mathbf{r} = (r_1 \cdots r_M)$ lie in the set $r_i \in \{0 \cdots N_{\mathbf{p}}\}$, with $r_i = 0$ if measurement i is due to clutter, and $r_i = j \neq 0$ if measurement i is generated by point feature j of the extended object.

In most practical applications the association hypothesis λ is unknown. We will nevertheless first derive a likelihood model for the measurements, conditional on a known association hypothesis. We will then show how the likelihood can be marginalised over the unknown hypotheses to remove this uncertainty. This strategy is only feasible if the total number of hypotheses is not too large. Later, in Section V, we will present an alternative strategy based on importance sampling to estimate the unknown association hypothesis.

Conditional on the association hypothesis we assume the measurements to be independent, so that the likelihood can be written as

$$p(\mathbf{y}|\mathbf{x}, \lambda) = \prod_{i=1}^M p(\mathbf{y}_i|\mathbf{x}, r_i), \quad (7)$$

with

$$p(\mathbf{y}_i|\mathbf{x}, r_i = j) = \begin{cases} \mathcal{U}_{\mathcal{Y}}(\mathbf{y}_i) = V^{-1} & \text{for } j = 0 \\ p(\mathbf{y}_i|\mathbf{p}_j^R) & \text{for } j = 1 \cdots N_{\mathbf{p}}, \end{cases} \quad (8)$$

where $\mathcal{U}_{\mathcal{A}}(\cdot)$ denotes the uniform distribution over the set \mathcal{A} . Thus, clutter measurements are assumed to be uniformly distributed over the range of the sensor \mathcal{Y} , a region whose volume is assumed to be V . For measurements of the point features the likelihood depends only on the location of the corresponding point feature in the fixed reference

¹In what follows we will suppress the time index for the sake of notational clarity.

frame. Under these assumptions the conditional likelihood can be written as

$$p(\mathbf{y}|\mathbf{x}, \lambda) = V^{-M_C} \prod_{i \in \mathcal{I}(\lambda)} p(\mathbf{y}_i | \mathbf{p}_{r_i}^R), \quad (9)$$

where $\mathcal{I}(\lambda)$ is the subset of measurement indices that correspond to point features. Note that the size of this set is M_T .

Example 2: If the sensors yield line of sight measurements of point sources relative to the observer, the individual measurements can be written as $\mathbf{y}_i = (R_i, \phi_i, \psi_i)$, where R_i is the range from the observer to the point source, and ϕ_i and ψ_i are the azimuth and elevation angles, respectively, of the point source relative to the observer. If the components of the line of sight measurements are assumed to be corrupted by independent Gaussian noise, the conditional likelihood for the point features becomes

$$p(\mathbf{y}_i | \mathbf{p}_j^R) = \mathcal{N}(\mathbf{y}_i | \hat{\mathbf{y}}(\mathbf{p}_j^R, \mathbf{p}^O), \mathbf{\Lambda}_y), \quad (10)$$

where we have assumed that $r_i = j \neq 0$. The covariance $\mathbf{\Lambda}_y = \text{diag}(\sigma_R^2, \sigma_\phi^2, \sigma_\psi^2)$ is the diagonal matrix with the individual measurement noise variances, and is assumed to be fixed and known. The components of the mean $\hat{\mathbf{y}}(\mathbf{p}_j^R, \mathbf{p}^O) = (\hat{R}_j, \hat{\phi}_j, \hat{\psi}_j)$ are given by

$$\begin{aligned} \hat{R}_j &= ((x_j^R - x^O)^2 + (y_j^R - y^O)^2 + (z_j^R - z^O)^2)^{1/2} \\ \hat{\phi}_j &= \tan^{-1} \left(\frac{y_j^R - y^O}{x_j^R - x^O} \right) \\ \hat{\psi}_j &= \tan^{-1} \left(\frac{z_j^R - z^O}{((x_j^R - x^O)^2 + (y_j^R - y^O)^2)^{1/2}} \right). \end{aligned} \quad (11)$$

For this model the range of the measurements is given by $\mathcal{Y} = [0, R_{\max}] \times [-\pi, \pi] \times [-\pi/2, \pi/2]$, where R_{\max} is the maximum range of the sensor. Thus, the volume of the measurement space becomes $V = 2\pi^2 R_{\max}$. ■

In most practical applications the association hypothesis λ is unknown. For a given number of point feature detections M_T in M measurements the total number of hypotheses are given by

$$N_\lambda(M_T, M_C) = \frac{M! N_{\mathbf{p}}!}{M_T! (M - M_T)! (N_{\mathbf{p}} - M_T)!}. \quad (12)$$

This follows from the number of ways of choosing a subset of M_T elements from the available M measurements, *i.e.* $\binom{M}{M_T}$, multiplied by the number of possible associations between the M_T detections and the $N_{\mathbf{p}}$ point features, *i.e.* $N_{\mathbf{p}}! / (N_{\mathbf{p}} - M_T)!$. However, the number of detected point features M_T is unknown, so that the total number of feasible hypotheses is given by

$$N_\lambda = \sum_{M_T=0}^{\min(N_{\mathbf{p}}, M)} N_\lambda(M_T, M_C). \quad (13)$$

If the total number of hypotheses is manageable, it is possible to marginalise over the association uncertainty, as is done in [39]. The marginal likelihood is of the form

$$p(\mathbf{y}|\mathbf{x}) = \sum_{\lambda} p(\lambda|\mathbf{x}) p(\mathbf{y}|\mathbf{x}, \lambda). \quad (14)$$

The marginalisation requires the definition of a prior distribution over the association hypothesis. We assume this prior to be of the form

$$p(\lambda|\mathbf{x}) = p(\mathbf{r}|M_T, M_C)p(M_C)p(M_T|\mathbf{x}), \quad (15)$$

with

$$\begin{aligned} p(\mathbf{r}|M_T, M_C) &= (N_\lambda(M_T, M_C))^{-1} \\ p(M_C) &= (\mu V)^{M_C} \exp(-\mu V)/M_C! \\ p(M_T|\mathbf{x}) &= \sum_h \prod_{i \in \mathcal{I}(h)} P_{D,i} \prod_{i \in \{1 \dots N_p\} - \mathcal{I}(h)} (1 - P_{D,i}). \end{aligned} \quad (16)$$

In the absence of measurements the prior for the association vector is assumed to be uniform over all the possible hypotheses for a known number of target detections. The number of clutter measurements is assumed to follow a Poisson distribution, with μ the spatial density of the clutter, which is assumed to be fixed and known. To compute the prior probability for the number of point feature detections we associate a unique detection probability with each of the point features $\{P_{D,i}\}_{i=1}^{N_p}$. Note that the detection probabilities may depend on the state, to model the effect that feature points may occlude each other in certain configurations relative to the observer. It is also possible to allow these probabilities to evolve over time, but we do not consider this scenario here. The prior probability for the number of detections is then obtained by summing over the $\binom{N_p}{M_T}$ possible ways to group M_T detections among the N_p point features, and h above ranges over these hypotheses. If the detection probability is the same for all the point features P_D , the prior for the number of point feature detections reduces to the binomial distribution, *i.e.*

$$p(M_T|\mathbf{x}) = \binom{N_p}{M_T} P_D^{M_T} (1 - P_D)^{N_p - M_T}. \quad (17)$$

Under these assumptions the marginal likelihood can be written as

$$p(\mathbf{y}|\mathbf{x}) \propto \sum_\lambda \left(\frac{P_D}{\mu(1 - P_D)} \right)^{M_T(\lambda)} \prod_{i \in \mathcal{I}(\lambda)} p(\mathbf{y}_i|\mathbf{p}_{r_i}^R), \quad (18)$$

where constant terms have been discarded, and we have made the dependency of the number of detections on the association hypothesis explicit, *i.e.* $M_T(\lambda)$.

Evaluation of the conditional likelihood in (9) is generally of $O(M)$ complexity. In contrast, the computational complexity for the marginal likelihood in (18) is $O(MN_\lambda)$. For a fixed number of point features and detection probability the number of hypotheses increases exponentially with an increase in the number of clutter measurements, as is depicted in Figure 2. Thus, in many applications of practical interest the number of feasible hypotheses is prohibitively large, so that the marginalisation cannot be performed explicitly. In Section V we will present an alternative importance sampling based strategy to take account of the association uncertainty.

IV. PARTICLE FILTER TRACKING

In this section we describe the particle filter tracking framework for a generic model parameterised by a state \mathbf{x} . We will present the specifics for the extended object tracking problem at the end of this section and in Section V.

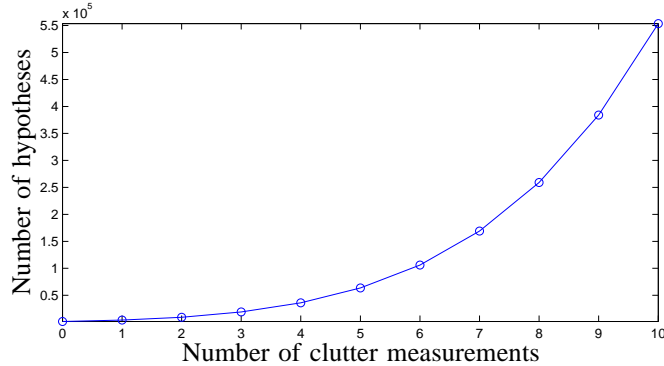


Fig. 2. **Complexity curve.** Number of hypotheses versus number of clutter measurements for five sources and a unity detection probability.

For tracking the distribution of interest is the posterior $p(\mathbf{x}_k|\mathbf{y}_{1:k})$, also known as the filtering distribution, where $\mathbf{y}_{1:k} = (\mathbf{y}_1 \cdots \mathbf{y}_k)$ denotes all the observations up to the current time step. In the Bayesian Sequential Estimation framework the filtering distribution can be computed according to the recursion

$$p(\mathbf{x}_k|\mathbf{y}_{1:k}) \propto p(\mathbf{y}_k|\mathbf{x}_k) \int p(\mathbf{x}_k|\mathbf{x}_{k-1})p(\mathbf{x}_{k-1}|\mathbf{y}_{1:k-1})d\mathbf{x}_{k-1}, \quad (19)$$

where $p(\mathbf{x}_k|\mathbf{x}_{k-1})$ is the dynamic model describing the state evolution, and $p(\mathbf{y}_k|\mathbf{x}_k)$ is the likelihood model. The recursion is initialised with some distribution for the initial state $p(\mathbf{x}_0)$. Once the sequence of filtering distributions is known point estimates of the state can be obtained according to any appropriate loss function, leading to, *e.g.*, Maximum *A Posteriori* (MAP) and Minimum Mean Square Error (MMSE) estimates.

The tracking recursion yields closed-form expressions in only a small number of cases. The most well-known of these is the Kalman filter [14] for linear Gaussian dynamic and likelihood models. Both the dynamic and likelihood models we consider here contain non-linear and non-Gaussian elements. This renders the tracking recursion analytically intractable, and approximation techniques are required. Sequential Monte Carlo (SMC) methods [19]–[21], otherwise known as Particle Filters [17], [18], or CONDENSATION [22], have gained a lot of popularity in recent years as a numerical approximation strategy to compute the tracking recursion for complex models. This is due to their simplicity, flexibility, ease of implementation, and modelling success over a wide range of challenging applications.

The basic idea behind particle filters is very simple. Starting with a weighted set of samples $\{\mathbf{x}_{k-1}^{(i)}, w_{k-1}^{(i)}\}_{i=1}^N$ approximately distributed according to $p(\mathbf{x}_{k-1}|\mathbf{y}_{1:k-1})$, new samples are generated from a suitably chosen proposal distribution, which may depend on the old state and the new measurements, *i.e.*, $\mathbf{x}_k^{(i)} \sim q(\mathbf{x}_k|\mathbf{x}_{k-1}^{(i)}, \mathbf{y}_k)$, $i = 1 \cdots N$. To maintain a consistent sample the new importance weights are set to

$$w_k^{(i)} \propto \frac{w_{k-1}^{(i)}p(\mathbf{y}_k|\mathbf{x}_k^{(i)})p(\mathbf{x}_k^{(i)}|\mathbf{x}_{k-1}^{(i)})}{q(\mathbf{x}_k^{(i)}|\mathbf{x}_{k-1}^{(i)}, \mathbf{y}_k)}, \quad \text{with} \quad \sum_{i=1}^N w_k^{(i)} = 1. \quad (20)$$

The new particle set $\{\mathbf{x}_k^{(i)}, w_k^{(i)}\}_{i=1}^N$ is then approximately distributed according to $p(\mathbf{x}_k|\mathbf{y}_{1:k})$. Approximations to the desired point estimates can then be obtained by Monte Carlo techniques. From time to time it is necessary

to resample the particles to avoid degeneracy of the importance weights. The resampling procedure essentially multiplies particles with high importance weights, and discards those with low importance weights. A full discussion of degeneracy and resampling falls outside the scope of this paper, but more details can be found in [19].

For the extended object tracking problem we will consider two scenarios. In the first setting we will assume that it is possible to marginalise over the association uncertainty, so that the dynamic and likelihood models in (3) and (18), respectively, apply. Similar to the bootstrap filter [17] we will take the proposal for the state to be the dynamics, so that the new importance weights become proportional to the corresponding particle likelihoods, multiplied by the old importance weights. The computational complexity of this algorithm is $O(NMN_\lambda)$ at each time step, and quickly becomes infeasible for realistic numbers of point features and clutter measurements. The second setting is described in the following section, and avoids the computationally expensive marginalisation by sampling the association hypotheses from an efficient proposal distribution.

V. TRACKING WITH ASSOCIATION UNCERTAINTY

In problems where the total number of association hypotheses is large direct marginalisation over the association uncertainty becomes computationally prohibitive. In this section we present an importance sampling based strategy to account for this uncertainty. More specifically, we augment the state with the unknown association hypothesis. Our aim will then be to estimate the posterior $p(\mathbf{x}_k, \lambda_k | \mathbf{y}_{1:k})$ recursively within the particle filtering framework. Under these assumptions the particle weights in (20) become

$$w_k^{(i)} \propto \frac{w_{k-1}^{(i)} p(\mathbf{y}_k | \mathbf{x}_k^{(i)}, \lambda_k^{(i)}) p(\lambda_k^{(i)} | \mathbf{x}_k^{(i)}) p(\mathbf{x}_k^{(i)} | \mathbf{x}_{k-1}^{(i)})}{q(\mathbf{x}_k^{(i)}, \lambda_k^{(i)} | \mathbf{x}_{k-1}^{(i)}, \lambda_{k-1}^{(i)}, \mathbf{y}_k)}, \quad (21)$$

where $p(\mathbf{y}_k | \mathbf{x}_k, \lambda_k)$ is the conditional likelihood in (9), $p(\lambda_k | \mathbf{x}_k)$ is the association hypothesis prior in (15), and $p(\mathbf{x}_k | \mathbf{x}_{k-1})$ is the dynamics in (3). The problem thus reduces to defining the joint proposal distribution for the state and association hypothesis $q(\mathbf{x}_k, \lambda_k | \mathbf{x}_{k-1}, \lambda_{k-1}, \mathbf{y}_k)$. We assume this proposal to be of the form

$$q(\mathbf{x}_k, \lambda_k | \mathbf{x}_{k-1}, \lambda_{k-1}, \mathbf{y}_k) = q(\lambda_k | \mathbf{x}_k, \mathbf{y}_k) p(\mathbf{x}_k | \mathbf{x}_{k-1}). \quad (22)$$

As before the proposal for the state is taken to be the dynamics in (3). The proposal for the association hypothesis depends only on the information available at the current time step. This reflects the fact that there is, at best, only a weak temporal dependence between the association hypotheses. We will define this proposal in terms of a proposal for the association vector $q(\mathbf{r}_k | \mathbf{x}_k, \mathbf{y}_k)$, with the proposals for the number of clutter measurements M_C and point feature detections M_T being implicit. As in [1], [33] we assume the proposal for the association vector to take the following factorised form²

$$q(\mathbf{r} | \mathbf{x}, \mathbf{y}) = \prod_{i=1}^M q(r_i | r_1 \cdots r_{i-1}, \mathbf{x}, \mathbf{y}_i). \quad (23)$$

²In what follows we will again suppress the time index for the sake of notational clarity.

Thus, the proposal for the i -th component of the association vector depends only on the corresponding measurement. Using Bayes' rule we define the proposal for the i -th component as

$$q(r_i|r_1 \cdots r_{i-1}, \mathbf{x}, \mathbf{y}_i) \propto p(\mathbf{y}_i|\mathbf{x}, r_i)q(r_i|r_1 \cdots r_{i-1}, \mathbf{x}). \quad (24)$$

The first term on the right hand side follows from the corresponding component of the conditional likelihood in (8), and of course will cancel with the corresponding numerator term in (21). For a particular point feature this value will be higher for measurements closer to the point feature, so that the proposal essentially captures the notion of a soft gating function. The second term forms the prior for the i -th component of the association vector, and we define it as

$$q(r_i = j|r_1 \cdots r_{i-1}, \mathbf{x}) \propto \begin{cases} q_0 & \text{for } j = 0 \\ q_j \prod_{k=1}^{i-1} (1 - \delta_{r_k, j}) & \text{for } j = 1 \cdots N_{\mathbf{p}}, \end{cases} \quad (25)$$

where $\delta_{i,j}$ denotes the Kronecker delta. The prior for the clutter hypothesis is set to be proportional to some fixed value $0 < q_0 < 1$. For all the other hypotheses the prior component is set to be proportional to some fixed value q_j , provided that the corresponding point feature has not already been assigned to another component of the association vector. For these hypotheses we will normally set the fixed value to be proportional to the detection probability of the corresponding point feature, *i.e.*

$$q_j = \frac{P_{D,j}(1 - q_0)}{\sum_{k=1}^{N_{\mathbf{p}}} P_{D,k}}. \quad (26)$$

This completes the specification of the proposal distribution for the i -th component of the association vector in (24). Since this distribution is discrete it is easily normalised and sampled from using standard techniques. Generating a sample for the entire association vector from the proposal in (23) can be achieved by sequentially sampling the individual components conditional on each other from r_1 to r_M . Note that the factorisation in (23) can be performed over any permutation of the components of the association vector. In practice we choose the order randomly for each particle at each iteration.

The computational complexity of the resulting particle filtering algorithm is $O(NMN_{\mathbf{p}})$. It increases only linearly with the number of point features and clutter measurements, and not exponentially, as is the case for the algorithm employing the marginal likelihood. This importance sampling approach to estimate the distribution of the association vector is more general than the application considered here. It can be applied in any setting, static or dynamic, where the data association problem arises within a probabilistic modelling framework.

VI. EXPERIMENTS AND RESULTS

In this section we report the results of two sets of experiments. The first, in Section VI-A, gives a comprehensive evaluation of the proposed extended object model and tracking algorithm on a challenging 3D trajectory under various filtering and modelling scenarios. The second, in section in Section VI-B, illustrates the sampling based data association strategy on synthetic 1D and 3D problems and compares with nearest-neighbour Kalman algorithms.

A. Experiment 1: Extended Object Tracking

For the extended object we will use the stick model from Example 1 in Section II, with the detection probability fixed to $P_D = 0.75$ for each of the point features. To evaluate the tracking performance we generated a deterministic trajectory with $K = 51$ unevenly spaced time steps, which we will denote by $\{\mathbf{x}_k^*\}_{k=1}^K$. For the observations we will take line of sight measurements from a non-stationary observer location, as in Example 2 in Section III. We set the maximum range for the sensor to $R_{\max} = 2,000$, and the measurement noise standard deviations to $(\sigma_R, \sigma_\phi, \sigma_\psi) = (2, 1^\circ, 1^\circ)$. The trajectory for the observer was also generated deterministically.

We consider two main scenarios. In the first there are no clutter measurements, whereas the second is characterised by an increasing level of clutter. In both cases we generate synthetic measurements by first sampling from the prior for the association hypothesis in (15), and conditional on this, sampling from the conditional likelihood model in (9).

For each of the scenarios we ran the particle filter for three different algorithm settings. In the first, which we will refer to as *truth*, we used the conditional likelihood in (9), conditional on the true hypothesis. This establishes a baseline performance for the other algorithms, as it is unrealisable in practice without prior knowledge of the association hypothesis. In the second, which we will refer to as *marginal*, we used the marginal likelihood in (18). In the third, which we will refer to as *conditional*, we again used the conditional likelihood in (9), but this time generated the association hypotheses from the proposal distribution in (23). In all cases the particle filter was initialised from a Gaussian distribution centred on the true first state, and state samples were generated from the dynamics in (3), with $(\sigma_x, \sigma_\beta, \sigma_\gamma, \sigma_p) = (1, 0.5^\circ, 1^\circ, 0)$. For the association proposal the prior for the clutter hypothesis was set to $q_0 = 10^{-6}$ for the experiments with clutter, and $q_0 = 0$ for those without. Using the particles, we computed MMSE estimates for the states as in (28). We compared the performance of the different algorithms in terms of the Root Mean Square Error (RMSE), which we define as

$$RMSE = \sqrt{\frac{1}{KN_{\mathbf{P}}} \sum_{k=1}^K \sum_{i=1}^{N_{\mathbf{P}}} \|\hat{\mathbf{p}}_{i,k}^R - \mathbf{p}_{i,k}^{R*}\|^2}, \quad (27)$$

and the computational effort. As a measure for the computational effort we took the average execution time per time step of a fairly optimal Matlab implementation of the different algorithms.

The results for the experiments without clutter are depicted in Figure 3. These were obtained by averaging over 20 runs of the algorithm for an increasing number of particles. For all the runs the same target trajectory was used, but new measurements were generated for each individual run. As expected the error decreases with an increase in the number of particles, with the performance of the truth algorithm being generally the best. The marginal algorithm outperforms the conditional algorithm for small numbers of particles. This is due to the larger state-space for the latter algorithm, making the search problem more complex. For more than 500 particles the error statistics for the three algorithms become virtually indistinguishable. For all three algorithms the computational effort increases linearly with an increase in the number of particles, with the conditional algorithm being slightly superior to the marginal algorithm. Thus, in the absence of clutter measurements, and for a small number of point features, the

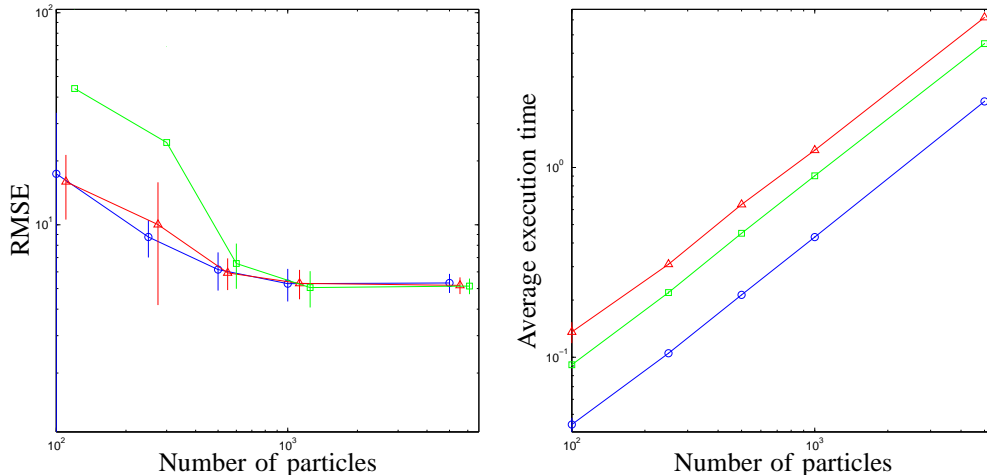


Fig. 3. **RMSE and average execution time statistics for experiments without clutter.** RMSE (left) and average execution time (right) for the true (blue circles), marginal (red triangles) and conditional (green squares) algorithms, as a function of the number of particles. The error decreases with an increase in the number of particles, with the performance of the three algorithms being virtually indistinguishable for 500 particles and above. For all three algorithms the computational effort increases linearly with the number of particles.

number of hypotheses is manageable.

The results for the experiments with clutter are depicted in Figure 4. These were obtained by averaging over 20 runs of the algorithm for an increasing clutter rate, with the number of particles fixed to $N = 500$. Again, all runs shared the same trajectory, with new measurements generated for each individual run. The error statistics are similar for each of the algorithms, and remain reasonably constant with an increase in the clutter rate. In this case, however, the computational effort increases exponentially with the clutter rate for the marginal algorithm. The computational effort is much lower for the conditional algorithm, and increases only slightly with an increasing clutter rate. To conclude Figure 5 shows the true and typical estimated trajectories for the first feature point for a clutter density of $\mu = 2/V$. The subjective quality of the trajectories is comparable for the three algorithms, but that for the marginal algorithm comes at a substantial computational cost.

B. Experiment 2: Comparison of Sampling Based Data Association with Kalman filter

We performed two simulations to illustrate the performance of sampling based data association. In both cases we assumed simple rigid target motion, and Gaussian Cartesian measurements around the true target locations. We assumed Gaussian random walk models for the target velocities; thus the unknown state vector for the model consists of the center of gravity for the target, its velocity, and the $N_p - 1$ relative positions of object features from the center of gravity. This was implemented using a Rao-Blackwellized particle filter [19], [20], [40] since we have assumed a rigid target and linear dynamics/observation equation. The idea of Rao-Blackwellization is simultaneous use of Kalman filter and particle filters. The Kalman filter is used for marginalisation of the continuous variable part of the state vector, and the particle filter is used for the discrete variable part of it, i.e. the association vector.

The first simulation considered 1D tracking of $N_p = 10$ target feature points over 30 time steps, with $P_D = 0.75$




Fig. 4. **RMSE and average execution time statistics for experiments with clutter.** RMSE (left) and average execution time (right) for the true (blue circles), marginal (red triangles) and conditional (green squares) algorithms, as a function of the clutter density, with the number of particles fixed to $N = 500$. The error statistics are similar for all three algorithms, and remain reasonably constant with an increase in the clutter rate. The computational effort for the marginal algorithm grows exponentially with the clutter rate, whereas that for the conditional algorithm increases only slightly.




Fig. 5. **Estimated trajectories.** True (blue circles) and estimated (red triangles) trajectories for the first point feature, obtained by the true algorithm (left), marginal algorithm (middle) and conditional algorithm (right). The trajectory of the observer is shown in black squares. The subjective quality is similar for all three algorithms.

and measurement noise standard deviation $\sigma = 0.1$. The second simulation considered a more challenging 3D tracking scenario of $N_p = 5$ target points over 50 time steps, with $P_D = 0.75$ and measurement noise standard deviation $\sigma = 0.3$. Variance parameters of the Gaussian random walk models for the target velocities are $\sigma_x = 0.1$ for the 1D scenario, and $\sigma_x = \sigma_y = \sigma_z = 1$ for the 3D scenario. We used $N = 1,000$ particles for both the 1D and the 3D scenarios.

Here, we explain how the tracker was initialised. Note that the focus of this paper is the evaluation of tracking performance, not on track initialisation. For this reason, the assumption on the initial state is rather well-conditioned. We suppose that first two observations contain neither missing observations nor clutter. The initial state of the tracker is then set as follows. For the center of gravity of targets in the initial state, a Gaussian distribution with mean at the center of gravity of observation at $k = 1$ and variance 1.0 is assumed. For the velocity of the initial state, a Gaussian

distribution is assumed with mean calculated by the difference between the centers of gravity of observations $k = 1$ and $k = 2$, and variance 0.1. For the remaining part of the state vector, which are the $N_p - 1$ relative positions of the target features from the center of gravity, a Gaussian distribution with mean 0 and variance 1.0 is used.

Synthetic data for the 1D scenario in an observation space of size 15, with clutter density $\mu = 1.0$ (i.e. an expected number of 15 clutter measurements per time slice) is shown in Figure 6. Typical estimation results for the 1D scenario are given in Figure 7. The plotted estimate is the Minimum Mean-Square Error estimate of the target state, obtained from the particle filter in the standard way as:

$$\hat{\mathbf{x}}_k = \sum_{i=1}^N w_k^{(i)} \mathbf{x}_k^{(i)}, \quad (28)$$

and this is seen to follow closely follow the true target trajectories.

We now conduct comparison experiments between the proposed particle filter method and 1) the Kalman filter given the true association, and 2) the nearest neighbor association Kalman filter. In the latter, the nearest observation is assigned to the corresponding target in state vector, with any observations out of a 4 sigma region being assigned automatically to clutter. Typical results for these two filters, operating on the same dataset as above, are shown in Figures 8 and 9. The result from the 1) (Kalman filter with true associations) is considered to be the optimal estimate since the true associations are given for the estimation and the remainder of the model is linear/Gaussian. Comparing with the result for the nearest neighbor association Kalman filter, the tracking is lost from around $k = 15$, owing to the high levels of clutter. More sophisticated data association could of course be adopted here, but we choose this basic scheme to give a baseline for comparison.

The three methods are now compared in Monte Carlo trials over a number of different scenarios. For clutter densities $\mu = 0.0, 0.2, 0.4, 0.6, 0.8, 1.0$ and 2.0 , we have observed the change of the performance for these three methods. Figure 10 shows comparison of the methods by root mean square errors in observation space. Figure 11 shows thee comparisons between Kalman filter given true associations and the proposed method. We can see in these figures that, for $\mu \geq 0.8$, the performance of the nearest neighbour association Kalman filter deteriorates significantly, while the deterioration of the proposed method is small compared to the optimal Kalman filter, as shown in Figure 11. These deteriorations compared to the optimal are mainly caused by incorrect association estimates within each method. To analyse this effect, we have summarized the number of wrong associations as shown in Figure 12. We can observe the significant difference in the numbers of wrong associations between the nearest neighbor association Kalman filter and the proposed method, especially for $\mu \geq 0.8$.

Synthetic data for the 3D scenario is shown in Figure 13. Estimation results by nearest neighbour association Kalman filter and the proposed method are shown in Figures 14 and 15. We can see in the figures that the result by the proposed method is close (almost indistinguishable from) the true trajectory, while by nearest neighbor association Kalman filter there are some noticeable discrepancies between the truth and the estimated.




Fig. 6. **Example Synthetic data for the 1D scenario.** Plotted points are the measurements. Target features are detected with probability $P_D = 0.75$, and the measurement noise standard deviation is $\sigma = 0.1$. Clutter density is $\mu = 1.0$. This is a challenging scenario which has many missed detections and high clutter level.




Fig. 7. **Estimation results for the 1D scenario.** Result for tracking of 10 points of an object in motion in 1D space. Lines show the estimated trajectories of the point features in measurement space and their true trajectories (the two are almost indistinguishable in this case). Measurements are shown as dots, and many clutter measurements are present with its density being $\mu = 1.0$.

VII. CONCLUSIONS

In this paper we considered the extended object tracking problem. Extended objects were modelled as point features in a target reference frame. We developed a realistic dynamic model to capture the motion of the extended object in terms of the translation and rotation of the target reference frame relative to a fixed reference frame. For the generally unlabelled measurements we showed how a multi-hypothesis likelihood can be obtained by marginalising over all the association hypotheses, as in [39]. The complexity of this marginalisation increases exponentially with an increase in the number of point features and clutter measurements, and quickly becomes infeasible in realistic




Fig. 8. **Estimation results for the 1D scenario by Kalman filter given the true association vector.** Result for tracking of 10 points on an object in motion in 1D space. Lines show the estimated trajectories of the point features in measurement space and their true trajectories. Measurements are denoted by dots, and many clutter measurements are present with density $\mu = 1.0$. Good results are obtained since the true association is given to the filter.

scenarios. As an alternative we proposed a particle filtering algorithm where the unknown association hypotheses were sampled from an efficiently designed proposal distribution. The computational complexity of this algorithm is substantially lower than that for an equivalent strategy using the marginal likelihood. Initial results show the estimation accuracy for the two strategies to be comparable. It should also be noted that the importance sampling approach to estimate the distribution of the association vector is more general than the application considered here. It can be applied to other tracking problems involving multiple targets, or in any setting, static or dynamic, where the data association problem arises within a probabilistic modelling framework.

One of the main considerations for future work is the design of a more efficient proposal for the state. Sampling the state from the dynamics takes no account of the new measurements, and it is often necessary to artificially increase the excitation noise of the dynamic models to ensure that all the viable regions of the state-space are explored with a finite number of particles. The downside of this is that the resulting trajectories may not be as smooth or accurate as desired. With proposals that take account of the new measurements the excitation parameters can be left at realistic levels, leading to smoother trajectories. Further important issues to be considered include the development of strategies to learn and adapt the parameters of the dynamic model online, for example in the spirit of [41], and to automatically detect the number of point features.

VIII. ACKNOWLEDGEMENTS

The work of Vermaak and Godsill was funded by QinetiQ under the project ‘Extended and Joint Object Tracking and Identification’, CU006-14890. The work of Ikoma was carried out under the ISM Cooperative Research Program (2004-ISM-CPR-2020), and was partially supported by the Ministry of Education, Science, Sports and Culture, Japan, Grant-in-Aid for Scientific Research (A), 14208025, 2004.




Fig. 9. **Estimation results for the 1D scenario by nearest neighbor association Kalman filter.** Result for tracking of 10 points on an object in motion in 1D space. Lines show the estimated trajectories of the point features in measurement space and their true trajectories. Measurements are denoted by dots, and many clutter measurements are present with density $\mu = 1.0$. The filter loses track from about $k = 15$.




Fig. 10. **Comparison of Root Mean Square Error of tracking result for the 1D scenario.** Root Mean Square Error of tracking result for 10 points on an object in motion in 1D space. KF (left bar) denotes the result by Kalman filter with true associations, NNKF (center bar) denotes the result by the nearest neighbor association Kalman filter, and RBPF (right hand bar) denotes that of by the proposed particle filter method.

REFERENCES

- [1] J. Vermaak, N. Ikoma, and S. J. Godsill, "Extended object tracking using particle techniques," in *Proceedings of the IEEE Aerospace Conference*, 2004.
- [2] D. Reid, "An algorithm for tracking multiple targets," *IEEE Transactions on Automation and Control*, vol. 24, no. 6, pp. 84–90, 1979.
- [3] Y. Bar-Shalom and T. E. Fortmann, *Tracking and Data Association*. Academic Press, 1988.
- [4] 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. 8, pp. 173–184, 1983.
- [5] L. Y. Pao, "Multisensor multitarget mixture reduction algorithms for target tracking," *AIAA Journal of Guidance, Control and Dynamics*, vol. 17, pp. 1205–1211, 1994.




Fig. 11. Comparison of Root Mean Square Error of tracking result for the 1D scenario between Kalman filter and the proposed method. Root Mean Square Error tracking result for 10 points on an object in motion in 1D space. KF (left bar) denotes the result by Kalman filter, and RBPF (right bar) denotes that of by the proposed method.




Fig. 12. Comparison of the number of wrong associations in tracking results for the 1D scenario. NNKF (left bar) denotes the result by the nearest neighbor association Kalman filter, and RBPF (right bar) denotes that of the proposed method.

- [6] D. J. Salmond, "Mixture reduction algorithms for target tracking in clutter," in *Signal and Data Processing of Small Targets, SPIE 1305*, O. E. Drummond, Ed., 1990, pp. 434–445.
- [7] H. Gauvrit, J.-P. L. Cadre, and C. Jauffret, "A formulation of multitarget tracking as an incomplete data problem," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 33, no. 4, pp. 1242–1257, 1997.
- [8] R. L. Streit and T. E. Luginbuhl, "Maximum likelihood method for probabilistic multi-hypothesis tracking," in *Signal and Data Processing of Small Targets, SPIE 2235*, O. E. Drummond, Ed., 1994.
- [9] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *Journal of the Royal Statistical Society, Series B*, vol. 39, no. 1, pp. 1–38, 1977.
- [10] P. Willett, Y. Ruan, and R. Streit, "PMHT: Problems and some solutions," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 38, no. 3, pp. 738–754, 2002.




Fig. 13. **Synthetic data for the 3D scenario.** Plots are the observations of the targets. Target features are detected with detection probability $P_D = 0.75$, and the measurement noise standard deviation being $\sigma = 0.3$.

- [11] J. R. Vasquez and J. L. Williams, "Improved hypothesis selection for multiple hypothesis tracking," in *Signal and Data Processing of Small Targets, SPIE 5428*, 2004.
- [12] J. L. Williams, "Gaussian mixture reduction for tracking multiple manoeuvring targets in clutter," Master's thesis, Air Force Institute of Technology, Wright-Patterson Air Force Base, 2003.
- [13] J. L. Williams and P. S. Maybeck, "Cost-function-based hypothesis control techniques for multiple hypothesis tracking," in *Signal and Data Processing of Small Targets, SPIE 5428*, 2004.
- [14] B. D. O. Anderson and J. B. Moore, *Optimal Filtering*. Englewood Cliffs: Prentice-Hall, 1979.
- [15] D. L. Alspach and H. W. Sorenson, "Nonlinear Bayesian estimation using Gaussian sum approximation," *IEEE Transactions on Automatic Control*, vol. 17, no. 4, pp. 439–448, 1972.
- [16] G. Kitagawa, "Non-Gaussian state-space modeling of nonstationary time series (with discussion)," *Journal of the American Statistical Association*, vol. 82, no. 400, pp. 1032–1063, 1987.
- [17] N. J. Gordon, D. J. Salmund, and A. F. M. Smith, "Novel approach to nonlinear/non-Gaussian Bayesian state estimation," *IEE Proceedings-F*, vol. 140, no. 2, pp. 107–113, 1993.
- [18] G. Kitagawa, "Monte Carlo filter and smoother for non-Gaussian nonlinear state space models," *Journal of Computational and Graphical Statistics*, vol. 5, no. 1, pp. 1–25, 1996.
- [19] A. Doucet, S. J. Godsill, and C. Andrieu, "On sequential Monte Carlo sampling methods for Bayesian filtering," *Statistics and Computing*, vol. 10, no. 3, pp. 197–208, 2000.
- [20] A. Doucet, J. F. G. de Freitas, and N. J. Gordon, Eds., *Sequential Monte Carlo Methods in Practice*. New York: Springer-Verlag, 2001.
- [21] J. S. Liu and R. Chen, "Sequential Monte Carlo methods for dynamic systems," *Journal of the American Statistical Association*, vol. 93, pp. 1032–1044, 1998.
- [22] 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.
- [23] F. Daum and J. Huang, "Nonlinear filtering with quasi-Monte Carlo methods," in *Signal and Data Processing of Small Targets, SPIE 5204*, 2003.
- [24] M. S. Oh, "Monte Carlo integration via importance sampling: Dimensionality effect and an adaptive algorithm," *Contemporary Mathematics*, vol. 115, pp. 165–187, 1991.
- [25] R. Karlsson and F. Gustafsson, "Monte Carlo data association for multiple target tracking," in *Proceedings of the IEE Seminar – Target Tracking: Algorithms and Applications*, 2001, pp. 13/1–13/5.




Fig. 14. **Estimation results for the 3D scenario by nearest neighbor association Kalman filter.** Result for tracking of 5 points on an object in motion in 3D space. Blue Lines show the estimated trajectories of the point features in measurement space and their true trajectories (black lines). Measurements are denoted by dots, and many clutter measurements were present. Some discrepancies between the true trajectory and the estimated one are observed.

- [26] D. Schulz, W. Burgard, D. Fox, and A. B. Cremers, "Tracking multiple moving targets with a mobile robot using particle filters and statistical data association," in *Proceedings of the IEEE International Conference on Robotics and Automation*, 2001, pp. 1665–1670.
- [27] C. Hue, J.-P. L. Cadre, and P. Pérez, "Tracking multiple objects with particle filtering," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 38, no. 3, pp. 791–812, 2002.
- [28] S. Geman and D. Geman, "Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. PAMI-6, no. 6, pp. 721–741, 1984.
- [29] F. Dellaert, S. M. Seitz, C. Thorpe, and S. Thrun, "EM, MCMC, and chain flipping for structure from motion with unknown correspondence," *Machine Learning*, vol. 50, no. 1-2, pp. 45–71, 2003.
- [30] W. R. Gilks, S. Richardson, and D. J. Spiegelhalter, *Markov Chain Monte Carlo in Practice*. Chapman and Hall, 1996.
- [31] N. J. Gordon and A. Doucet, "Sequential Monte Carlo for manoeuvring target tracking in clutter," in *Signal and Data Processing of Small Targets, SPIE 3809*, O. E. Drummond, Ed., 1999, pp. 493–500.
- [32] A. Doucet, N. J. Gordon, and V. Krishnamurthy, "Particle filters for state estimation of jump Markov linear systems," *IEEE Transactions on Signal Processing*, vol. 49, no. 3, pp. 613–624, 2001.
- [33] N. Ikoma and S. J. Godsill, "Extended object tracking with unknown association, missing observations, and clutter using particle filters," in *Proceedings of the 12th IEEE Workshop on Statistical Signal Processing*, 2003, pp. 485–488.
- [34] S. J. Godsill and J. Vermaak, "Models and algorithms for tracking using trans-dimensional sequential Monte Carlo," in *Proceedings of the IEEE International Conference on Acoustic, Speech and Signal Processing*, 2004.
- [35] J. Vermaak, C. Andrieu, A. Doucet, and S. J. Godsill, "Particle methods for Bayesian modeling and enhancement of speech signals," *IEEE Transactions on Speech and Audio Processing*, vol. 10, no. 3, pp. 173–185, 2002.
- [36] J. Vermaak, A. Doucet, and P. Pérez, "Maintaining multi-modality through mixture tracking," in *Proceedings of the IEEE International Conference on Computer Vision*, 2003, pp. II: 1110–1116.
- [37] J. Vermaak, S. J. Godsill, and A. Doucet, "Sequential Bayesian kernel regression," in *Advances in Neural Information Processing Systems 16*, S. Thrun, L. Saul, and B. Schölkopf, Eds., 2004.
- [38] —, "Radial basis function regression using trans-dimensional sequential Monte Carlo," in *Proceedings of the 12th IEEE Workshop on Statistical Signal Processing*, 2003, pp. 525–528.
- [39] N. J. Gordon, D. J. Salmond, and D. Fisher, "Bayesian target tracking after group pattern distortion," in *Signal and Data Processing of*




Fig. 15. **Estimation results for the 3D scenario by the proposed method.** Result for tracking of 5 points on an object in motion in 3D space. Blue lines show the estimated trajectories of the point features in measurement space and their true trajectories (black lines - almost indistinguishable from blue).

Small Targets, SPIE 3163, O. E. Drummond, Ed., 1997, pp. 238–248.

[40] G. Casella and C. P. Robert, “Rao-blackwellisation of sampling schemes,” *Biometrika*, vol. 83, no. 1, pp. 81–94, 1996.

[41] G. Kitagawa, “A self-organizing state-space model,” *Journal of the American Statistical Association*, vol. 93, no. 443, pp. 1203–1215, 1998.