VARIABLE RATE PARTICLE FILTERS FOR TRACKING APPLICATIONS

Simon Godsill and Jaco Vermaak

Signal Processing Group
Department of Engineering, University of Cambridge, U.K.
sjg@eng.cam.ac.uk, jv211@eng.cam.ac.uk

ABSTRACT

Here we describe recent advances in particle filtering algorithms and models for tracking of manoeuvring objects in clutter. The methods develop on the basic variable dimension particle filtering algorithms introduced in [1], in which a new type of dynamical model is introduced whose state variables arrive at unknown times relative to the observation process (hence 'variable rate'). Targets are assumed to follow deterministic trajectories in between state times, determined by an appropriate model, such as the differential equation model for the object. The framework allows for automatic modelling and estimation of the trajectories of targets using an adaptation of particle filtering methods [2] into the variable dimension setting. In this paper we introduce more effective sampling schemes for the variable rate setting that ensure future states are only generated as and when required, new dynamical models appropriate for manoeuvring objects, and new observation models under the assumption of a non-homogeneous Poisson process for both targets and clutter. Simulations show very effective tracking performance under challenging settings which cannot be emulated in a standard fixed rate scheme.

1. INTRODUCTION

Filtering methods for tracking of objects in noise are now very well established, see for example [3, 4]. In straightforward settings (close to linear/Gaussian models, low clutter levels) the classical Kalman filter and its variants can be successfully adopted. In more challenging settings (highly non-linear, non-Gaussian models, high clutter densities and low detection probabilities), other numerical methods are required, and in recent years a class of methods receiving great attention is sequential Monte Carlo, or particle filters [2, 5, 6]. In common with the earlier methods the particle filter relies on a state-space representation of the system in terms of Markov hidden states $\{x_n\}$ and observations $\{y_n\}$:

$$\begin{aligned} x_{n+1} &\sim f(x_{n+1}|x_n) && \text{State evolution density} \\ y_n &\sim g(y_n|x_n) && \text{Observation density} \end{aligned} \qquad (1)$$

with $x_0 \sim f(x_0)$ being the initialisation, where $\sim$ denotes that the variable to the left is drawn independently from the probability density on the right.

The optimal ('Bayesian') filtering recursions from time index n to $n+1$ are then given by

$$p(x_{n+1}|y_{0:n}) = \int p(x_n|y_{0:n}) f(x_{n+1}|x_n) dx_n \qquad (2)$$

$$p(x_{n+1}|y_{0:n+1}) = \frac{g(y_{n+1}|x_{n+1}) p(x_{n+1}|y_{0:n})}{p(y_{n+1}|y_{0:n})} \qquad (3)$$

The Kalman filter implements this exactly in the linear/Gaussian setting. The general update rule is analytically intractable for most models of practical interest. We therefore turn to Sequential Monte Carlo (SMC) methods [2, 5, 6], also known as particle filters, to provide an efficient numerical approximation to the update rule. These methods have gained tremendous popularity in recent years over a wide range of tracking applications. They are applicable to non-linear and non-Gaussian models, and are able to capture multimodal distributions. The basic idea behind particle filters is simple: the target distribution is represented by a weighted set of Monte Carlo samples. These samples are propagated and updated using a sequential version of importance sampling as new measurements become available. Using the samples, estimates of the target state can be obtained using standard Monte Carlo integration techniques. The particle filter implements the filtering recursions approximately by propagation of a weighted 'cloud' of N particles $\{x_n^{(i)}, w_n^{(i)}, i = 1, ..., N\}$, $\sum_i w_n^{(i)} = 1$. A basic filter [2, 7] implements the following recursion at time n, and for particles $i = 1, ..., N$:

$$x_{n+1}^{(i)} \sim q(x_{n+1}|x_n^{(i)}, y_{0:n+1}),$$

$$w_{n+1}^{(i)} \propto w_n^{(i)} \frac{g(y_{n+1}|x_{n+1}^{(i)}) f(x_{n+1}^{(i)}|x_n^{(i)})}{q(x_{n+1}^{(i)}|x_n^{(i)}, y_{0:n+1})}$$

where $q()$ is a specially designed proposal density, and this is usually followed at some time steps by a resampling operation that selects the particles according to their weights and then resets weights to $1/N$.

There are now many applications of particle filters in the tracking arena, see e.g. [8]. Most of these applications apply particle filtering to fairly simple (typically linear) dynamical models with more or less elaborate observation models, including sometimes multiple targets and/or random clutter. In this paper we develop strategies based upon a variable rate state arrival process. The material takes as its basis the variable dimension particle filters and modelling framework introduced in [1], in which target states are allowed to arrive at different and unknown rates compared with the observation process. In this way the models are able to model parsimoniously the various turning/straight, smooth/non-smooth manoeuvres of an object. The basic particle filtering techniques required and basic tracking models were proposed in [1], and here we develop these further in several ways, by incorporation of: more efficient sampling schemes for the state process; specialised deterministic resampling methods; new variable rate dynamical models within an intrinsic coordinate system for the motion of manoeuvring targets; observation models are proposed based upon a Poisson assumption for both targets and clutter (in this way the data association problem in standard tracking work is avoided altogether and large numbers of clutter points/low target detection probabilities may readily be dealt with). We believe the new Poisson observation models adopted are closely related to those proposed in [9], although the derivation and exposition here are more straightforward. Simulations are presented for challenging scenarios with many clutter points and highly manoeuvrable targets, demonstrating the robustness of the variable rate filters compared with their more standard fixed rate counterparts.

2. VARIABLE RATE MODEL

The variable rate state is defined as as $\mathbf{x}_k = (\boldsymbol{\theta}_k, \tau_k)$, where $k \in \mathbb{N}$ is the discrete state index, $\tau_k \in \mathbb{R}^+$ denotes the state arrival time, and $\boldsymbol{\theta}_k$ denotes the vector of variables necessary to parameterise the target state (see [1]). Commonly they will include position, velocity, heading, *etc.*, variables. We will assume that the variable rate state sequence follows a Markovian process such that successive states are independently drawn as follows

$$\mathbf{x}_k \sim f(\mathbf{x}_k|\mathbf{x}_{k-1}) = f_{\boldsymbol{\theta}}(\boldsymbol{\theta}_k|\boldsymbol{\theta}_{k-1}, \tau_k, \tau_{k-1}) f_\tau(\tau_k|\tau_{k-1}). \tag{4}$$

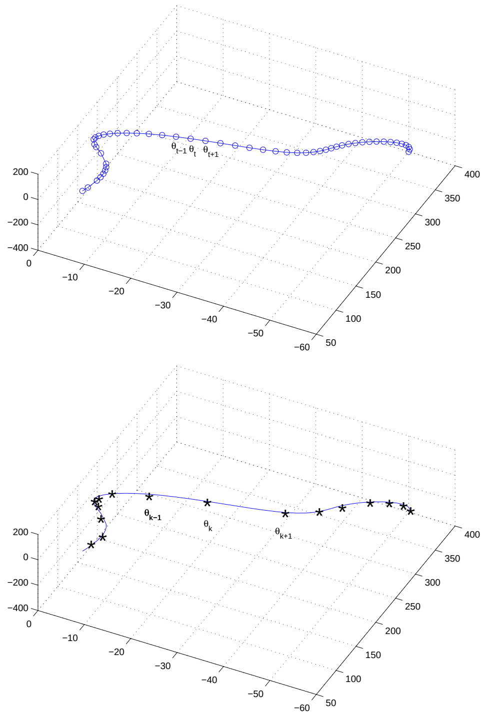

Fig. 1. Standard fixed rate representation (top) and variable rate representation (bottom)

The Markovian assumption is convenient for presentation, but can easily be relaxed if the models require it. See Fig. 2 for a comparison of the variable rate model with the fixed rate model. In the fixed rate model, trajectories are parameterised in terms of a state variable per time point, even when a very smooth and unvarying manoeuvre is being executed. The variable rate model, through a suitable choice of dynamical model, parameterises smooth straight sections with only a few state points, while rapid manoeuveres require more closely spaced states. We will denote by $\mathbf{y}_n$ the vector of measurements, with $n \in \mathbb{N}$ the discrete measurement time index, as above. The rate of the measurement process will typically (but not necessarily) be higher than that of the state process. In our variable rate setting the observation functions are specified in terms of a 'neighbourhood' $\mathcal{N}_n$ of state indices defining the dependence structure of the data. The data are then assumed drawn independently from a density function conditional on the neighbourhood of state points $\mathcal{N}_n$:

$$\mathbf{y}_n \sim g(\mathbf{y}_n|\{\mathbf{x}_k : k \in \mathcal{N}_n\}) = g(\mathbf{y}_n|\mathbf{x}_{\mathcal{N}_n}). \tag{5}$$

Note that the neighbourhood is a deterministic function of the time n and state sequence $\mathbf{x}_{1:\ldots}$. This is quite a general formulation that encompasses many

Bayesian models. Typically the likelihood is calculated as $g(\mathbf{y}_n|\hat{\theta}_n(\mathbf{x}_{\mathcal{N}_n}))$, where $\hat{\theta}_n(\mathbf{x}_{\mathcal{N}_n})$ is a deterministic function of $\mathbf{x}_{\mathcal{N}_n}$. We will propose specific examples of both dynamical models and observation functions in a later section.

3. VARIABLE RATE STATE ESTIMATION

The objective will be to estimate recursively the sequence of variable rate state points as the measurements become available. All the information concerning the variable rate state sequence is contained in its conditional probability distribution. In keeping with the standard filtering nomenclature this conditional distribution will be referred to as the variable rate filtering distribution, defined as

$$p(\mathbf{x}_{1:k_n^+}|\mathbf{y}_{1:n}), \tag{6}$$

where $\mathbf{y}_{1:n} = (\mathbf{y}_1 \cdots \mathbf{y}_n)$ is the sequence of available measurements up to time index n and $\mathbf{x}_{1:k} = (\mathbf{x}_1 \cdots \mathbf{x}_k)$ denotes a sequence of k hidden state variables; k_n^+ denotes the index of the state variable having largest time index τ_k within all neighbourhoods $\mathcal{N}_1$ through $\mathcal{N}_n$, i.e.

$$k_n^+ = \{k \in \mathcal{N}_{1:n} : \tau_k > \tau_l, \forall l \in \mathcal{N}_{1:n}, l \neq k\} \tag{7}$$

where $\mathcal{N}_{1:n} = \cup_{k=1}^n \mathcal{N}_k$. Note that k_n^+ is a random variable that depends deterministically on the sequence of x_k values.

The variable rate filtering distribution has variable dimension support since k_n^+ itself is a random variable (to be estimated along with the hidden state sequence). Hence the problem can be considered as one of dynamical model uncertainty as well as parameter estimation.

For a recursive inference procedure we require an update rule of the form

$$p(\mathbf{x}_{1:k_{n-1}^+}|\mathbf{y}_{1:n-1}) \xrightarrow{\mathbf{y}_n} p(\mathbf{x}_{1:k_n^+}|\mathbf{y}_{1:n}), \tag{8}$$

i.e. once a new measurement is received, the variable rate filtering distribution at the previous time step is incrementally updated to yield the new variable rate filtering distribution at the current time step. Using Bayes' theorem and the modelling assumptions of the variable rate model, the new variable rate filtering distribution can be related to that at the previous time step by

$$\begin{aligned} &p(\mathbf{x}_{1:k_n^+}|\mathbf{y}_{1:n}) \\ &\quad \propto g(\mathbf{y}_n|\mathbf{x}_{\mathcal{N}_n}) \frac{p(\mathbf{y}_{1:n-1}|\mathbf{x}_{1:k_n^+})}{p(\mathbf{y}_{1:n-1}|\mathbf{x}_{1:k_{n-1}^+})} \\ &\qquad f(\mathbf{x}_{k_{n-1}^+ +1:k_n^+}|\mathbf{x}_{k_{n-1}^+}) p(\mathbf{x}_{1:k_{n-1}^+}|\mathbf{y}_{1:n-1}). \end{aligned}$$

The first term in the above expression is the observation likelihood, whereas the third term denotes a repeated application of the variable rate state evolution model in (4) to complete the local neighbourhood for the new measurement $\mathbf{y}_n$, defined as

$$f(\mathbf{x}_{k:k+L}|\mathbf{x}_{k-1}) = \prod_{l=k}^{k+L} f(\mathbf{x}_l|\mathbf{x}_{l-1}). \tag{9}$$

It is worthwhile to note that the second term, the ratio of likelihoods, will be unity in many cases, since the addition of states $\mathbf{x}_{k_{n-1}^+ +1:k_n^+}$ will not alter the neighbourhood structures for data points $\mathbf{y}_{1:n-1}$ and hence the ratio of likelihoods is one. In practice this can be ensured for all of the models of interest here by a suitable definition of the neighbourhood structure. We will assume in all future calculations that neighbourhoods are constructed in this way so that future states beyond k_{n-1}^+ cannot alter the neighbourhood structure for $\mathbf{y}_{n-1}$, and the update equation then simplifies to

$$\begin{aligned} &p(\mathbf{x}_{1:k_n^+}|\mathbf{y}_{1:n}) \\ &\quad \propto g(\mathbf{y}_n|\mathbf{x}_{\mathcal{N}_n}) f(\mathbf{x}_{k_{n-1}^+ +1:k_n^+}|\mathbf{x}_{k_{n-1}^+}) p(\mathbf{x}_{1:k_{n-1}^+}|\mathbf{y}_{1:n-1}). \end{aligned}$$

This sequential update rule is similar in form to that for standard state-space models, comprising a likelihood and a dynamic component. There is, however, a very important difference: owing to the form of the variable rate model the number of state points required to represent any section of the target trajectory is an unknown random variable. This will play an important role in the development of efficient numerical techniques to implement the update.

The basic form of the particle filter required for this task is as in [1], summarised here for convenience. Assuming that we have a set of weighted samples approximately distributed according to the variable rate filtering distribution at the previous time step, *i.e.*

$$\{\mathbf{x}_{1:k_{n-1}^+}^{(i)}, w_{n-1}^{(i)}\}_{i=1}^N \sim p(\mathbf{x}_{1:k_{n-1}^+}|\mathbf{y}_{1:n-1}),$$

the particle filter update step proceeds as below when a new measurement becomes available, for each $i = 1, ..., N$:

$$\begin{aligned} &\mathbf{x}_{k_{n-1}^+ +1:k_n^+}^{(i)} \sim q(\mathbf{x}_{k_{n-1}^+ +1:k_n^+}|\mathbf{x}_{k_{n-1}^+}^{(i)}, \mathbf{y}_n) \\ &\qquad w_n^{(i)} \propto w_{n-1}^{(i)} \frac{g(\mathbf{y}_n|\mathbf{x}_{\mathcal{N}_n}^{(i)}) f(\mathbf{x}_{k_{n-1}^+ +1:k_n^+}^{(i)}|\mathbf{x}_{k_{n-1}^+}^{(i)})}{q(\mathbf{x}_{k_{n-1}^+ +1:k_n^+}^{(i)}|\mathbf{x}_{k_{n-1}^+}^{(i)}, \mathbf{y}_n)}, \end{aligned}$$

and as before resampling may be carried out as and when required, resetting the weights to $1/N$. $q()$ is any appropriate proposal function which can sequentially generate new state points conditional on previous ones. Note that the proposal mechanism involves repeatedly increasing the number of state points until the neighbourhood for time n is complete, i.e. the observation likelihood can be evaluated. In the examples here, we set $q()$ equal to the dynamical model for the variable rate states, i.e. (9).

The resulting set of weighted samples will then be approximately distributed according to the new variable rate filtering distribution, *i.e.*

$$\{\mathbf{x}_{1:k_n^+}^{(i)}, w_n^{(i)}\}_{i=1}^N \sim p(\mathbf{x}_{1:k_n^+}|\mathbf{y}_{1:n}).$$

In this paper we present two modifications to the basic algorithm above:

Deterministic resampling. In order to retain greater diversity in the particle population, particles are deterministically replicated at each time step (see [6] for general discussion of this issue). Specifically, if a particle has weight w_n greater greater than $1/N$, then we assign a multiplicity $m = \lfloor N w_n \rfloor$ to that particle, and a weight w_n/m. If a particle has weight w_n less than $1/N$, we preserve it in the particle set with a multiplicity of 1, and retaining its own weight w_n. This procedure preserves exactly the same representation of the filtering density, but maintains some lower weight particles that would normally be resampled out. Of course, such a procedure can lead to (at most) a doubling of the number of particles at each time step, and so a further pruning stage is included that preserves only the N highest weighted particles in the filter following the update to the next time step.

State regeneration. This addresses an inefficiency in the basic filter, since state times τ_k, $k \in \mathcal{N}_n$ are often well in the future (i.e. significantly greater than n). In the basic filter above these state points are never regenerated and can attain very low weights at some point in the future if they are not consistent with new data points. We address this by allowing for regeneration of these future state points at each time n. There are several correct ways to achieve this, but we have settled on one effective procedure. In essence we attempt to regenerate any part of the state that does not affect the current (or previous) likelihood functions $g(\mathbf{y}_n|\mathbf{x}_{\mathcal{N}_n})$. Typically this will involve regenerating one or more future τ_k values, conditional upon the sequence of values $\hat{\theta}_{1:n}$ used in likelihood computation. To see how this works, consider augmenting the variable rate state at time n with the value of $\hat{\theta}_n(\mathbf{x}_{\mathcal{N}_n})$, i.e. the state becomes $\{\mathbf{x}_{\mathcal{N}_n}, \hat{\theta}_n(\mathbf{x}_{\mathcal{N}_n})\}$. Now, suppose that some parts of $x_{\mathcal{N}_n}$ are 'redundant' in the sense that they can be modified without changing the values of $\hat{\theta}_{1:n}$. We can then resample these elements according to $p(\mathbf{x}_{\mathcal{N}_n}|\hat{\theta}_{1:n}, \mathbf{x}_{\mathcal{N}_{1:n-1}}, \mathbf{y}_{1:n}) = p(\mathbf{x}_{\mathcal{N}_n}|\hat{\theta}_{1:n}, \mathbf{x}_{\mathcal{N}_{1:n-1}})$

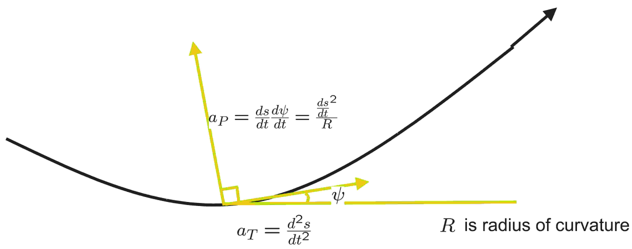

Fig. 2. Intrinsic coordinate system

(a degenerate distribution in general) and the particle weights remain unchanged as before, since this can be regarded as a very simple Markov chain transition kernel of the type described in [10]. Furthermore, since high weighted particles have a multiplicity greater than 1 (see above), some degree of stratification is often possible, since different particles can be drawn from different regions of the probability distribution.

4. VARIABLE RATE TRACKING MODELS

A new dynamical model appropriate for manoeuvring objects under the variable rate model is proposed, based on an intrinsic coordinate system. In an intrinsic coordinate system applied forces can be represented relative to the heading of the object, rather than relative to the more standard Cartesian or polar fixed coordinate frame. This we postulate is a more realistic representation of the thrusts applied when turning a vehicle. Distance travelled along the path of motion is denoted s, while angle of the path relative to horizontal is denoted ψ. Accelerations tangential to and perpendicular to the motion are then given in Fig. 2.In the variable rate model we now assume that a piecewise constant thrust, relative to the direction of heading, is applied between any two times τ_k and τ_{k+1}, with tangential component T_T and perpendicular component T_P. A resistance term $\lambda\frac{ds}{dt}$, assumed to apply in the opposite direction to the heading adds some damping to the system. Resolving forces tangentially and perpendicularly:

$$T_T = \lambda \frac{ds}{dt} + m\frac{d^2s}{dt^2}$$
$$T_P = m\frac{ds}{dt}\frac{d\psi}{dt}$$

These equations are readily integrated to give equations for $s(t)$ and $\psi(t)$ during any time period with fixed thrusts T_T and T_P, from which the cartesian position $x(t)$ can be obtained by numerical integration for any time t between τ_k and τ_{k+1}.

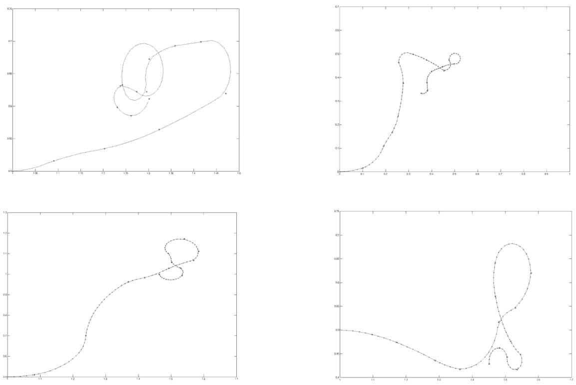

Fig. 3. Example trajectories simulated from the intrinsic coordinate model

We thus have the variable rate state variables as follows:

$$\theta_k = [T_{T,k},\, T_{P,k},\, v(\tau_k),\, \psi(\tau_k),\, x(\tau_k)]$$

from which the position, speed and angle at any time between τ_k and τ_{k+1} may be obtained deterministically. This implies that the variable rate neighbourhood structure for this model need contain only one element, i.e. $\mathcal{N}_t = \{k;\, \tau_k \leq t, \tau_{k+1} > t\}$.

To complete the dynamical model we specify the distribution of thrusts in the time interval τ_k to τ_{k+1}:

$$T_{T,k} \sim N(\mu_T, \sigma_T^2),\;\; T_{P,k} \sim N(0, \sigma_P^2)$$

and the distribution of time points:

$$\tau_{k+1} - \tau_k \sim G(\alpha_\tau, \beta_\tau)$$

where N is the normal and G is the gamma distribution.

Some example trajectories from the model are shown in Fig. 3, showing the clear ability of the model to generate elaborate manoeuvres.

Many possible observation models are available for this dynamical model, and we have explored several. In order to test a challenging non-Gaussian case, we adopt here a Poisson target model with independent Poisson clutter observations. The number of measurements from the target at each time point are assumed drawn randomly from a Poisson distribution having mean λ_T. Each such target measurement spatially then has a Gaussian distribution centered on the true position. Random clutter measurements are also included, having mean number λ_C and uniform spatial distribution. See [9] for a very similar model for extended object tracking. The model for the measurements is then:

$$y_{n,i} \sim \begin{cases} N(x_n, \sigma_y^2), & \text{target measurement} \\ U_R, & \text{clutter measurement} \end{cases}$$

where U_R is the uniform distribution over the region of surveillance R. Under these assumptions it can be shown from the properties of Poisson point processes that

$$g(\mathbf{y}_n | x_n) \propto \Pi_{i=1}^{N_n} \lambda_{yc}(y_{n,i}),$$
$$\lambda_{yc}(y) = \lambda_T N(y|\mathbf{x}_n, \sigma_y^2 I) + \lambda_C U_R$$

where N_n is the total number of measurements at time index n. This model thus avoids any explicit treatment of the data association problem inherent in many tracking scenarios, but can potentially handle very high clutter densities.

5. RESULTS

We have experimented successfully in various different scenarios. To demonstrate this in one representative case, see the following example. Here we have a surveillance region of size 2000×2000, clutter parameter $\lambda_C = 20$ (i.e. an average of 20 clutter points per time point) and target mean $\lambda_T = 0.7$, i.e. the target is often unobserved, and with observation parameter $\sigma_y = 40$. A manoeuvring target is generated using the intrinsic dynamic model with $m = 500$, $\lambda = 0.3$, $\mu_T = 50$, $\sigma_T = 50$ and $\sigma_P = 1000$. The data and results are summarised in Fig. 4. Note that the target appears almost buried in clutter, but the variable rate filter is able to track its manoeuvres successfully. A more detailed analysis shows that the filter is able to maintain many feasible hypotheses when the data is ambiguous. By contrast the corresponding fixed rate filter is unable to follow the trajectory and loses track. Although it has the same intrinsic model for the dynamics, it is unable to track the sustained turning manoeuvres from the variable rate model.

6. ACKNOWLEDGEMENTS

The authors acknowledge the suggestion of a Poisson observation model for extended targets by Kevin Gilholm and David Salmond, and helpful discussion with Simon Maskell, all at QinetiQ UK. The work of both authors was partially supported by QinetiQ Contract CU006 0000014890 under the Weapon, Platform and Effectors domain of the MOD UK research programme, and that of Godsill by the UK Government's Data and Information Fusion DTC.

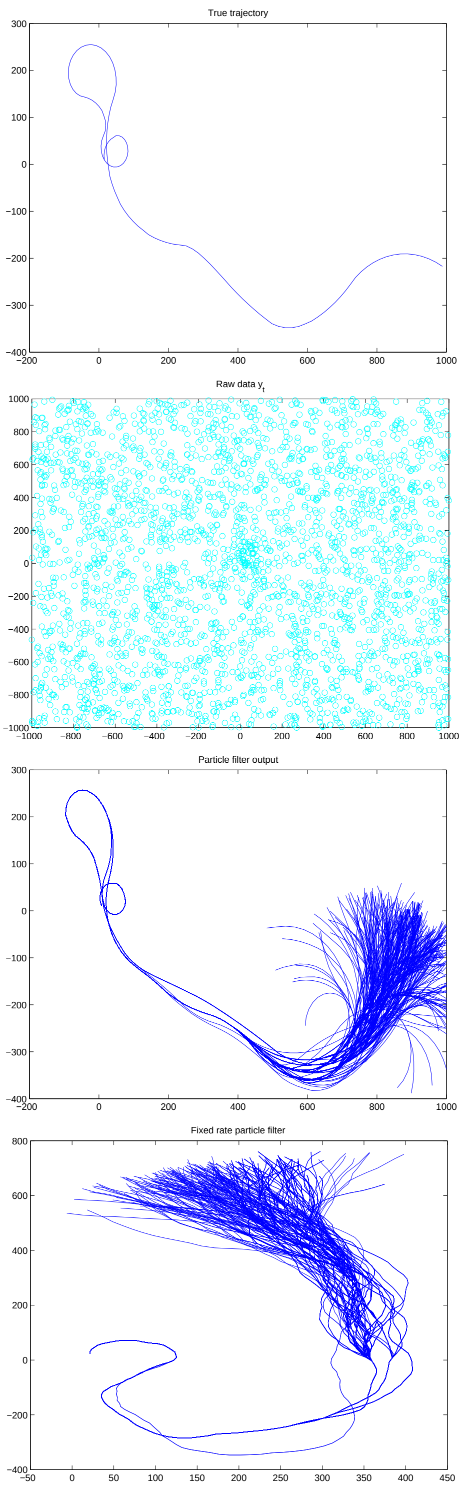

Fig. 4. Tracking example - Poisson likelihood, 400 particles. Top - true trajectory, second down - raw data, third down - variable rate particle filter (all particles shown), bottom - fixed rate particle filter (all particles shown)

7. REFERENCES

[1] S. J. Godsill and J. Vermaak, "Models and algorithms for tracking using trans-dimensional sequential Monte Carlo," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing*, 2004.

[2] A. Doucet, S. Godsill, and C. Andrieu, "On sequential Monte Carlo sampling methods for Bayesian filtering," *Statistics and Computing*, vol. 10, pp. 197–208, 2000.

[3] Y. Bar-Shalom and W. D. Blair, *Multitarget-Multisensor Tracking: Applications and Advances*, vol. III, Archtech House, Norwood, MS, 2000.

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

[5] N.J. Gordon, D.J. Salmond, and A.F.M. Smith, "Novel approach to non-linear/non-Gaussian Bayesian state estimation," *IEE Proceedings-F*, vol. 140, no. 2, pp. 107–113, 1993.

[6] A. Doucet, N. de Freitas, and N. Gordon, Eds., *Sequential Monte Carlo in Practice*, Springer-Verlag, New York, 2001.

[7] J. Liu and R. Chen, "Sequential Monte Carlo Methods for Dynamic Systems," *Journal of the American Statistical Association*, vol. 93, no. 443, pp. 1032–1044, 1993.

[8] B. Ristic, S. Arulampalam, and N. Gordon, *Beyond the Kalman Filter: Particle Filters for Tracking Applications*, Artech House Books, 2004.

[9] K. Gilholm and D. J. Salmond, "Extended object and group tracking," in *Symposium on Target Tracking and Sensor Data Fusion for Military Observation Systems*, 2003.

[10] S N MacEachern, M Clyde, and J S Liu, "Sequential importance sampling for nonparametric Bayes models: the next generation," *Can. J. Stats.*, vol. 27, pp. 251–267, 1999.