PREPRINT: IEEE TRANS. SIGNAL PROCESSING, VOL. 59, NO. 9, pp. 4473-4477, SEPT. 2011

Bernoulli Forward-Backward Smoothing for Joint Target Detection and Tracking

Ba-Tuong Vo*, Daniel Clark†, Ba-Ngu Vo‡, Branko Ristic§

Abstract

In this paper we derive a forward-backward smoother for joint target detection and estimation and propose a sequential Monte Carlo implementation. We model the target by a Bernoulli random finite set since the target can be in one of two 'present' or 'absent' modes. Finite Set Statistics is used to derive the smoothing recursion. Our results indicate that smoothing has two distinct advantages over just using filtering: Firstly, we are able to more accurately identify the appearance and disappearance of a target in the scene and secondly, we can provide improved state estimates when the target exists.

Index Terms

Tracking, filtering, detection, estimation, smoothing.

---

* Corresponding author. Address: School of Electrical, Electronic and Computer Engineering, The University of Western Australia, Crawley, WA 6009, Australia. E-mail: ba-tuong.vo@uwa.edu.au Tel: +61 8 6488 1767 Fax: +61 8 6488 1065

† Address: Joint Research Institute in Signal & Image Processing, School of Engineering & Physical Sciences Heriot-Watt University, Edinburgh, EH14 4AS, UK. E-mail: d.e.clark@hw.ac.uk Tel: +44(0) 131 451 3286 Fax: +44(0) 131 451 4155

‡ Address: School of Electrical, Electronic and Computer Engineering, The University of Western Australia, Crawley, WA 6009, Australia. E-mail: ba-ngu.vo@uwa.edu.au Tel: +61 8 6488 1767 Fax: +61 8 6488 1065

§ Address: ISR Division, Defence Science and Technology Organisation, Fishermans Bend, VIC 3207, Australia. E-mail: branko.ristic@dsto.defence.gov.au Tel: +61 3 9626 8226 Fax: +61 3 9626 8341

I. Introduction

There has been considerable interest in filtering and smoothing in the last half-century, motivated by the discovery in 1960 of the solution to the linear filtering problem [1]. Smoothing differs from prediction and filtering in that the estimate of the state of the system at a specific point in time is to be determined from a batch of measurements, some of which may be collected later than the time-step we are interested in. Consequently, there is a delay in producing the estimate of the state at that time, though more accurate estimates can be obtained since more information about the system is available. Since the 1960's, soon after the solutions to various filtering problems came corresponding solutions to the smoothing problems [2], [3]. More recent work on non-linear filtering and smoothing has been inspired by sequential Monte Carlo approximations to the Bayesian interpretation of filtering theory [4].

In surveillance applications, the target of interest may not always be present in the surveillance region. A target can enter and exit the surveillance region at random instances. Moreover, due to background clutter, exact knowledge of target existence in the surveillance area cannot be assumed. A filter that does not account for existence of targets may follow spurious measurements when the target is not in the scene and when the target enters the scene the tracker may not be able to lock-on to the target. Thus, it is crucial that the filter detects the presence of the target as well as tracking it. The Bernoulli filter or Joint Target-Detection and Tracking (JoTT) filter is a solution to the problem of joint detection and tracking of a single target under the presence of detection uncertainty and clutter. In this paper, we derive the forward-backward smoother for the Bernoulli model. Numerical results are presented using a sequential Monte Carlo implementation. Preliminary results have been announced in [5].

II. Joint target detection and estimation

In a joint detection estimation problem, it is assumed that at most one target can be present, and that the target can be in one of two 'present' or 'absent' modes. Thus, we model the target state X_k, at time k, as finite set, which can take on either the empty set or a singleton. Mahler's Finite Set Statistics (FISST) provides practical mathematical tools for dealing with finite-set-valued random variable [6], [7], including a consistent notion of integration and density. A *Bernoulli* random finite set (RFS) on $\mathcal{X}$, with parameters (r, p), has probability $1 - r$ of being empty, and probability r of being a singleton whose (only) element is distributed according to a probability density p (defined on $\mathcal{X}$). The FISST probability density of a Bernoulli RFS is (see [7] pp. 368)

$$\pi(X) = \begin{cases} 1 - r & X = \emptyset, \\ r \cdot p(x) & X = \{x\} \end{cases} . \tag{1}$$

We write the (FISST) density of a Bernoulli density as $\pi = (r, p)$.

A. Bernoulli state space model

For the dynamical model, let X_{k-1}, which can take on either the empty set $\emptyset$ or $\{x_{k-1}\}$, denote the state at time $k-1$. Conditional upon $X_{k-1}=\emptyset$, the target can re-enter the scene with probability $p_{R,k|k-1}$ and occupy kinematic state x_k with probability density $f_{R,k|k-1}(x_k)$, or remain absent from the scene with probability $1-p_{R,k|k-1}$. More concisely, conditional upon $X_{k-1}=\emptyset$, X_k is the Bernoulli RFS described by the density $\phi_{k|k-1}(\cdot|\emptyset)=(p_{R,k|k-1},f_{R,k|k-1})$, i.e.

$$\phi_{k|k-1}(X_k|\emptyset)=\begin{cases}1-p_{R,k|k-1}, & X_k=\emptyset,\\ p_{R,k|k-1}f_{R,k|k-1}(x_k), & X_k=\{x_k\}\end{cases}. \tag{2}$$

In addition, conditional upon $X_{k-1}=\{x_{k-1}\}$, with probability $p_{S,k|k-1}(x_{k-1})$ the target can survive to the next time step and take on a new state x_k with probability density $f_{k|k-1}(x_k|x_{k-1})$, or disappear with probability $1-p_{S,k|k-1}(x_{k-1})$. In other words, conditional upon $X_{k-1}=\{x_{k-1}\}$, X_k is the Bernoulli RFS described by the density $\phi_{k|k-1}(\cdot|\{x_{k-1}\})=(p_{S,k|k-1},f_{k|k-1}(\cdot|x_{k-1}))$, i.e.

$$\phi_{k|k-1}(X_k|\{x_{k-1}\})=\begin{cases}1-p_{S,k|k-1}(x_{k-1}), & X_k=\emptyset,\\ p_{S,k|k-1}(x_{k-1})f_{k|k-1}(x_k|x_{k-1}) & X_k=\{x_k\}\end{cases}. \tag{3}$$

For the measurement model, let X_k, which can take on either the empty set $\emptyset$ or $\{x_k\}$, denote the state at time k. Conditional upon $X_k=\{x_k\}$, the measurements follow the standard single target in clutter model [7], [8] with likelihood

$$\eta_k(Z_k|x)=e^{-\langle\kappa_k,1\rangle}\left([1-p_{D,k}(x)]\kappa_k^{Z_k}+p_{D,k}(x)\sum_{z\in Z_k}g_k(z|x)\kappa_k^{Z_k-\{z\}}\right) \tag{4}$$

where $p_{D,k}(x_k)$ is the probability of detection, $g_k(z|x_k)$ is the likelihood of the target generated measurement z, κ_k is the intensity of clutter, $\kappa_k^{Z_k}=\prod_{z\in Z_k}\kappa_k(z)$, and $Z_k-\{z\}$ denotes a set difference. On the other hand conditional upon $X_k=\emptyset$, all measurements must originate from clutter. Hence, the likelihood of the measurement set Z_k at time k is

$$\gamma_k(Z_k|X_k)=\begin{cases}e^{-\langle 1,\kappa_k\rangle}\kappa_k^{Z_k}, & X_k=\emptyset\\ \eta_k(Z_k|x) & X_k=\{x\}\end{cases}.$$

B. Bernoulli filter

Let $\pi_{k|k-1}(\cdot|Z_{1:k-1})$, $\pi_{k|k}(\cdot|Z_{1:k})$, denote the predicted and filtered multi-target densities. Then, using the FISST notion of integration and density, filtered density can be propagated forward in time by the multi-target Bayes filter [6], [7]

$$\pi_{k|k-1}(X_k|Z_{1:k-1})=\int\phi_{k|k-1}(X_k|X)\pi_{k-1|k-1}(X|Z_{1:k-1})\delta X, \tag{5}$$

$$\pi_{k|k}(X_k|Z_{1:k})=\frac{\gamma_k(Z_k|X_k)\pi_{k|k-1}(X_k|Z_{1:k-1})}{\int\gamma_k(Z_k|X)\pi_{k|k-1}(X|Z_{1:k-1})\delta X}, \tag{6}$$

where

$$\int f(X)\delta X = \sum_{i=0}^{\infty} \frac{1}{i!} \int f(\{x_1, ..., x_i\}) dx_1 \cdots dx_i.$$

is the set integral of a function f, which takes the finite subsets of the space $\mathcal{X}$ to the reals.

Under the Bernoulli model, an exact recursion for the predicted and filtering densities can easily be obtained as stated in the following proposition. This result can be verified by substituting the transition density and likelihood function into the multi-target Bayes filter [9].

Proposition 1 *If the initial prior density is a Bernoulli, then under the Bernoulli state space model all subsequent predicted and filtered densities are Bernoulli. Moreover let* $\pi_{k-1|k-1} = (r_{k-1|k-1}, p_{k-1|k-1})$ *denote the filtered density at time* $k-1$*, then the predicted density to time* k *is* $\pi_{k|k-1} = (r_{k|k-1}, p_{k|k-1})$*, where*

$$r_{k|k-1} = p_{R,k|k-1}(1 - r_{k-1|k-1}) + r_{k-1|k-1} \int p_{S,k|k-1}(x) p_{k-1|k-1}(x) dx \tag{7}$$

$$p_{k|k-1}(\zeta) = \frac{p_{R,k|k-1}(1 - r_{k-1|k-1}) f_{R,k|k-1}(\zeta)}{r_{k|k-1}} + \frac{r_{k-1|k-1} \int p_{S,k|k-1}(x) f_{k|k-1}(\zeta|x) p_{k-1|k-1}(x) dx}{r_{k|k-1}} \tag{8}$$

and the filtered density at time k *is* $\pi_{k|k} = (r_{k|k}, p_{k|k})$ *where*

$$r_{k|k} = \frac{r_{k|k-1} \int \eta_k(Z_k|x) p_{k|k-1}(x) dx}{(1 - r_{k|k-1}) e^{-\langle \kappa_k, 1 \rangle} \kappa_k^{Z_k} + r_{k|k-1} \int \eta_k(Z_k|x) p_{k|k-1}(x) dx}, \tag{9}$$

$$p_{k|k}(x) = \frac{\eta_k(Z_k|x) p_{k|k-1}(x)}{\int \eta_k(Z_k|x) p_{k|k-1}(x) dx}. \tag{10}$$

Notice that the propagation of the probability of target existence $r_{k|k}$ is now coupled to the propagation of the distribution $p_{k|k}$ of the kinematic state. The Bernoulli RFS filter used in the paper has been independently derived by Mahler as the JoTT filter in [7], section 14.7, and [9].

C. Forward-Backward Bernoulli Smoother

Forward-backward smoothing consists of forward filtering followed by backward smoothing. In the forward pass, the filtering density is propagated forward to time l via the Bayes recursion. In the backward pass, the smoothed density is propagated backward, from time l to time $k < l$, via the backward smoothing recursion (see for example [10]). In the set-valued case, the filtering density is propagated forward via the multi-target Bayes recursion (5), (6) and the smoothed density $\pi_{k|l}$ is propagated backward, via the multi-target backward smoothing recursion

$$\pi_{k-1|l}(X) = \pi_{k-1|k-1}(X) \int \phi_{k|k-1}(Y|X) \frac{\pi_{k|l}(Y)}{\pi_{k|k-1}(Y)} \delta Y. \tag{11}$$

The recursion (11) has the same form as the standard backward smoother expressed in terms of ordinary densities and integrals. A simple way to derive (11) is to first apply the same argument as per the standard backward smoother to relevant RFS probability densities, then invoke the relationship between FISST density/integration with measure theoretic density/integration in [11].

Proposition 2 *If the initial prior density is a Bernoulli, then under the Bernoulli state space model all smoothed densities are Bernoulli. Moreover let* $\pi_{k|l} = (r_{k|l}, p_{k|l})$ *denote the smoothed Bernoulli density from time* l *to* k*, then the smoothed Bernoulli density from time* l *to* $k-1$ *is* $\pi_{k-1|l} = (r_{k-1|l}, p_{k-1|l})$*, where*

$$r_{k-1|l} = 1 - (1 - r_{k-1|k-1}) \left(\alpha_{R,k|l} + \beta_{R,k|l} \int \frac{p_{k|l}(\zeta)}{p_{k|k-1}(\zeta)} f_{R,k|k-1}(\zeta) d\zeta \right) \tag{12}$$

$$p_{k-1|l}(x) = \frac{p_{k-1|k-1}(x) \left(\alpha_{S,k|l}(x) + \beta_{S,k|l}(x) \int \frac{p_{k|l}(\zeta)}{p_{k|k-1}(\zeta)} f_{k|k-1}(\zeta|x) d\zeta \right)}{\int p_{k-1|k-1}(x) \alpha_{S,k|l}(x) dx + \int p_{k-1|k-1}(x) \beta_{S,k|l}(x) \int \frac{p_{k|l}(\zeta)}{p_{k|k-1}(\zeta)} f_{k|k-1}(\zeta|x) d\zeta dx} \tag{13}$$

with

$$\alpha_{R,k|l} = (1 - p_{R,k|k-1}) \frac{(1 - r_{k|l})}{(1 - r_{k|k-1})}, \tag{14}$$

$$\beta_{R,k|l} = p_{R,k|k-1} \frac{r_{k|l}}{r_{k|k-1}}, \tag{15}$$

$$\alpha_{S,k|l}(x) = (1 - p_{S,k|k-1}(x)) \frac{(1 - r_{k|l})}{(1 - r_{k|k-1})}, \tag{16}$$

$$\beta_{S,k|l}(x) = p_{S,k|k-1}(x) \frac{r_{k|l}}{r_{k|k-1}}. \tag{17}$$

Proof: Suppose that $\pi_{k|l}$ is Bernoulli. Then, it follows from (11) that

$$\begin{aligned} \pi_{k-1|l}(X) &= \pi_{k-1|k-1}(X) \left(\phi_{k|k-1}(\emptyset|X) \frac{\pi_{k|l}(\emptyset)}{\pi_{k|k-1}(\emptyset)} + \int \phi_{k|k-1}(\{y\}|X) \frac{\pi_{k|l}(\{y\})}{\pi_{k|k-1}(\{y\})} dy \right) \\ &= \pi_{k-1|k-1}(X) \left(\phi_{k|k-1}(\emptyset|X) \left(\frac{1 - r_{k|l}}{1 - r_{k|k-1}} \right) + \frac{r_{k|l}}{r_{k|k-1}} \int \phi_{k|k-1}(\{y\}|X) \frac{p_{k|l}(y)}{p_{k|k-1}(y)} dy \right) \end{aligned}$$

For $X = \emptyset$, $\pi_{k-1|k-1}(X) = 1 - r_{k-1|k-1}$, $\phi_{k|k-1}(\emptyset|X) = 1 - p_{R,k|k-1}$, $\phi_{k|k-1}(\{y\}|X) = p_{R,k|k-1} f_{R,k|k-1}(y)$, hence using (14), (15) gives

$$\pi_{k-1|l}(\emptyset) = (1 - r_{k-1|k-1}) \left(\alpha_{R,k|l} + \beta_{R,k|l} \int f_{R,k|k-1}(y) \frac{p_{k|l}(y)}{p_{k|k-1}(y)} dy \right) = 1 - r_{k-1|l}$$

For $X = \{x\}$, $\pi_{k-1|k-1}(X) = r_{k-1|k-1} p_{k-1|k-1}(x)$, $\phi_{k|k-1}(\emptyset|X) = 1 - p_{S,k|k-1}(x)$, $\phi_{k|k-1}(\{y\}|X) = p_{S,k|k-1}(x) f_{k|k-1}(y|x)$, hence using (16), (17) gives

$$\pi_{k-1|l}(\{x\}) = r_{k-1|k-1} p_{k-1|k-1}(x) \left(\alpha_{S,k|l}(x) + \beta_{S,k|l}(x) \int f_{k|k-1}(y|x) \frac{p_{k|l}(y)}{p_{k|k-1}(y)} dy \right)$$

As a function of x, the above expression does not necessarily integrate to 1. Therefore it is necessary to normalize which then gives (13). The smoothed density $\pi_{k-1|l}$ is indeed a Bernoulli since from the property of a FISST density $\int \pi_{k-1|l}(X) \delta x = 1$, and as a consequence $r_{k-1|l} = \int \pi_{k-1|l}(\{x\}) dx$. This equation is not obvious from the expression for $\pi_{k-1|l}(\{x\})$ and $r_{k-1|l}$. However, this can be verified by noting that

$$\int \pi_{k-1|l}(\{x\})dx = r_{k-1|k-1}\int p_{k-1|k-1}(x)\alpha_{S,k|l}(x)dx + r_{k-1|k-1}\iint p_{k-1|k-1}(x)\beta_{S,k|l}(x)f_{k|k-1}(y|x)\frac{p_{k|l}(y)}{p_{k|k-1}(y)}dydx$$

For the first integral on the RHS of the above, substituting (16) for $\alpha_{S,k|l}(x)$, then rearrange using (7) and (14) gives

$$\int p_{k-1|k-1}(x)\alpha_{S,k|l}(x)dx = \frac{1-r_{k|l}-\alpha_{R,k|l}(1-r_{k-1|k-1})}{r_{k-1|k-1}}$$

For the double integral, substituting (17) for $\beta_{S,k|l}(x)$, then rearrange using (8) and (17) gives

$$\iint p_{k-1|k-1}(x)\beta_{S,k|l}(x)f_{k|k-1}(y|x)\frac{p_{k|l}(y)}{p_{k|k-1}(y)}dydx = \frac{r_{k|l}-(1-r_{k-1|k-1})\beta_{R,k|k-1}\int\frac{f_{R,k|k-1}(y)}{p_{k|k-1}(y)}p_{k|l}(y)dy}{r_{k-1|k-1}}$$

Hence, $\int \pi_{k-1|l}(\{x\})dx = r_{k-1|l}$. Since $\pi_{l|l}$ is Bernoulli by virtue of Proposition 1, it follows by induction that each smoothed density is also Bernoulli.

III. Sequential Monte Carlo Implementation

A. Prediction

The prediction involves a sum of two terms similar to the PHD filter [6] and the implementation is adopted from [11]. Suppose at time $k-1$ that the Bernoulli density $\pi_{k-1|k-1}$ is represented by $r_{k-1|k-1}$ and a set of weighted particles $\{w_{k-1}^{(i)}, x_{k-1}^{(i)}\}_{i=1}^{N_{k-1}}$, i.e.

$$p_{k-1|k-1}(x) \approx \sum_{i=1}^{N_{k-1}} w_{k-1}^{(i)}\delta_{x_{k-1}^{(i)}}(x).$$

Then the predicted probability of target existence and track density are approximated with

$$r_{k|k-1} \approx p_{R,k|k-1}(1-r_{k-1|k-1}) + r_{k-1|k-1}\sum_{i=1}^{N_{k-1}} w_{k-1}^{(i)} p_{S,k|k-1}(x_{k-1}^{(i)})$$

$$p_{k|k-1}(x) \approx \sum_{i=1}^{N_{k-1}+J_k} w_{k|k-1}^{(i)}\delta_{x_k^{(i)}}(x_k)$$

where the particles are drawn from two proposal distributions with appropriate support,

$$x_k^{(i)} \sim \begin{cases} q_k\left(\cdot|x_{k-1}^{(i)}, Z_k\right), & i = 1,...,N_{k-1} \\ s_k\left(\cdot|Z_k\right), & i = N_{k-1}+1,...,N_{k-1}+J_k \end{cases}$$

$$w_{k|k-1}^{(i)} = \begin{cases} \frac{r_{k-1|k-1}}{r_{k|k-1}}\frac{p_{S,k|k-1}(x_{k-1}^{(i)})f_{k|k-1}\left(x_k^{(i)}|x_{k-1}^{(i)}\right)w_{k-1}^{(i)}}{q_k\left(x_k^{(i)}|x_{k-1}^{(i)},Z_k\right)}, \\ \frac{1-r_{k-1|k-1}}{r_{k|k-1}}\frac{p_{R,k|k-1}f_{R,k|k-1}\left(x_k^{(i)}\right)}{J_k s_k\left(x_k^{(i)}|Z_k\right)}. \end{cases}$$

B. Update

After the prediction, suppose at time k that the predicted density $\pi_{k|k-1}$ is represented by $r_{k|k-1}$ and a set of weighted particles $\{w_{k|k-1}^{(i)}, x_k^{(i)}\}_{i=1}^{N_{k|k-1}}$, i.e.

$$\pi_{k|k-1}(x) \approx \sum_{i=1}^{N_{k|k-1}} w_{k|k-1}^{(i)} \delta_{x_k^{(i)}}(x).$$

Then, the updated probability of target existence is

$$r_{k|k} = \frac{r_{k|k-1} \sum_{i=1}^{N_{k|k-1}} w_{k|k-1}^{(i)} \eta_k(Z_k|x_{k|k-1}^{(i)})}{(1-r_{k|k-1})e^{-\langle \kappa_k, 1\rangle} \kappa_k^{Z_k} + r_{k|k-1} \sum_{i=1}^{N_{k|k-1}} w_{k|k-1}^{(i)} \eta_k(Z_k|x_{k|k-1}^{(i)})}$$

and the updated track density is

$$p_{k|k}(x) \approx \sum_{i=1}^{N_{k|k-1}} w_k^{(i)} \delta_{x_k^{(i)}}(x)$$

where

$$w_k^{(i)} = \tilde{w}_k^{(i)} / \sum\nolimits_{i=1}^{N_{k|k-1}} \tilde{w}_k^{(i)},$$

$$\tilde{w}_k^{(i)} = \eta_k(Z_k|x_{k|k-1}^{(i)}) w_{k|k-1}^{(i)}.$$

The recursion is initialized by generating a set of weighted particles $\{w_0^{(i)}, x_0^{(i)}\}_{i=1}^{N_0}$ representing $p_{0|0}$.

C. Backward Smoothing

The backward smoothing works backward from time l and uses the smoothed density at time l with the filtered density for each time-step to obtain a smoothed particle approximation of the single-target density and existence probability.

Assume we have $r_{k|k}$, $r_{k+1|l}$ and the particle approximations

$$p_{k|k}(x) \approx \sum_{i=1}^{N_{k|k}} w_k^{(i)} \delta_{x_k^{(i)}}(x)$$

$$p_{k+1|l}(x) \approx \sum_{j=1}^{N_{k+1|l}} w_{k+1|l}^{(j)} \delta_{x_{k+1|l}^{(j)}}(x)$$

Then, the particle approximations of $r_{k|l}$, and $p_{k|l}$ are given by

$$r_{k|l} \approx 1 - (1 - r_{k|k}) \left(\frac{(1 - r_{k+1|l})}{(1 - r_{k+1|k})} \left(1 - p_{R,k+1|k}\right) + \frac{r_{k+1|l}}{r_{k+1|k}} p_{R,k+1|k} \sum_{j=1}^{N_{k+1|l}} w_{k+1|l}^{(j)} \frac{f_{R,k+1|k}(x_{k+1|l}^{(j)})}{p_{k+1|k}(x_{k+1|l}^{(j)})} \right)$$

$$p_{k|l}(x) \approx \sum_{i=1}^{N_{k|k}} \tilde{w}_{k|l}^{(i)} \delta_{x_k^{(i)}}(x)$$

where

$$\tilde{w}_{k|l}^{(i)} = \frac{w_{k|l}^{(n)}}{\sum_{n=1}^{N_{k|k}} w_{k|l}^{(n)}}$$

$$w_{k|l}^{(i)} = \frac{(1 - r_{k+1|l})}{(1 - r_{k+1|k})} \left(1 - p_{S,k+1|k}(x_k^{(i)})\right) w_k^{(i)} + \frac{r_{k+1|l}}{r_{k+1|k}} \sum_{j=1}^{N_{k+1|l}} p_{S,k+1|k}(x_k^{(i)}) w_{k+1|l}^{(j)} \frac{f_{k+1|k}(x_{k+1|l}^{(j)}|x_k^{(i)})}{p_{k+1|k}(x_{k+1|l}^{(j)})} w_k^{(i)}$$

and the predicted density is computed by $p_{k+1|k}(x_{k+1|l}^{(j)}) = \sum_{i=1}^{N_{k|k}} w_k^{(i)} f_{k+1|k}(x_{k+1|l}^{(j)}|x_k^{(i)})$.

The convergence of the particle approximation to the integrals in equations (12) and (13) can be verified with well understood property of empirical measures (see chapters 1 and 2 of [4]).

IV. Numerical Studies

In this section, we present results of the proposed sequential Monte Carlo forward-backward smoother and illustrate the resulting performance improvement over filtering.

We consider a nearly constant turn model with varying turn rate together with bearing and range measurements. The observation region is the half disc of radius $2000m$. The state variable $x_k = [\tilde{x}_k^T \omega_k]^T$ comprises the planar position and velocity $\tilde{x}_k^T = [p_{x,k}, \dot{p}_{x,k}, p_{y,k}, \dot{p}_{y,k}]$ as well as the turn rate ω_k. The state transition model is

$$\tilde{x}_k = F(\omega_{k-1})\tilde{x}_{k-1} + G\nu_{k-1},$$
$$\omega_k = \omega_{k-1} + \Delta u_{k-1},$$

where

$$F(\omega) = \begin{bmatrix} 1 & \frac{\sin \omega\Delta}{\omega} & 0 & -\frac{1-\cos \omega\Delta}{\omega} \\ 0 & \cos \omega\Delta & 0 & -\sin \omega\Delta \\ 0 & 1 - \frac{\cos \omega\Delta}{\omega} & 1 & \frac{\sin \omega\Delta}{\omega} \\ 0 & \sin \omega\Delta & 0 & \cos \omega\Delta \end{bmatrix}.$$

$$G = \begin{bmatrix} \frac{\Delta^2}{2} & 0 \\ \Delta & 0 \\ 0 & \frac{\Delta^2}{2} \\ 0 & \Delta \end{bmatrix}.$$

$\Delta = 1s$, $\nu_{k-1} \sim \mathcal{N}(\cdot; 0, \sigma_\nu^2)$ with $\sigma_\nu = 5ms^{-2}$, and $u_{k-1} \sim \mathcal{N}(\cdot; 0, \sigma_u^2)$ with $\sigma_u = \pi/180rads^{-2}$. The probability of survival is $p_{S,k|k-1} = 0.99$. The observation region is the half disc $[-\pi/2, \pi/2]rad \times [0, 2000]m$. The probability of target birth is $p_{R,k|k-1} = 0.01$ and the birth state density is $f_{R,k|k-1}(x) = \mathcal{N}(x; m_{R,k|k-1}, Q_{R,k|k-1})$ where $m_{R,k|k-1} = [-11, -13, 107, 21, 0.1]$ and $Q_{R,k|k-1} = diag([100, 10, 100, 10, 0.01]^2)$. The target measurement is a noisy bearing and range vector

$$z_k = [\arctan(p_{x,k}/p_{y,k}), \sqrt{p_{x,k}^2 + p_{y,k}^2}]^T + \epsilon_k,$$

where $\epsilon_k \sim \mathcal{N}(\cdot; 0, R_k)$, with $R_k = diag([\sigma_\theta^2, \sigma_r^2])$, $\sigma_\theta = 1.5(\pi/180)rad$, and $\sigma_r = 10m$. The probability of detection is $p_{D,k} = 0.88$. Clutter follows a Poisson RFS with a mean rate of 30 returns per scan and a uniform spatial distribution on the observation region.

At each time step, $N = 1000$ are sampled from the prediction for the proposal, $J = 1000$ new particles are sampled directly from the birth density, and 1000 particles are resampled from the posterior following each update.

State estimation is performed with the following simple procedure. A target is declared present if the estimated probability of existence is greater than 0.5, otherwise no target is declared. If a target is declared, the state estimate is given by the mean of the posterior state distribution, otherwise if no target is declared, there is no state estimate.

We evaluate filter performance using the Optimal Sub-Pattern Assignment (OSPA) multi-target miss-distance. We use the OSPA metric because it jointly captures differences in cardinality and individual elements between two finite sets in a mathematically consistent yet intuitively meaningful way [12].

For joint detection and tracking we only need the OSPA distance between two finite sets with cardinality of at most one. The construction of the OSPA distance $d_{OSPA}^{(c)}(X, Y)$ between two finite sets X and Y with cardinality of at most one is as follows. $d_{OSPA}^{(c)}(\emptyset, \emptyset) = 0$, $d_{OSPA}^{(c)}(\{x\}, \{y\}) = \min(\|x - y\|, c)$, $d_{OSPA}^{(c)}(\{x\}, \emptyset) = d_{OSPA}^{(c)}(\emptyset, \{x\}) = c$. The cut-off parameter c determines the relative weighting of the penalties assigned to cardinality and localization errors. For further details see [12].

The simulation is run over 100 time steps, where a target was initialized at time step $k = 11$ and terminated at time step $k = 94$. The forward filter is run as per usual, and the backward smoother is run for a lag of 3 time steps. An example simulation is given in Figures 1 and 2, showing the measurements, true trajectory and the resulting filtered and smoothed estimates. It can be seen that the filter initiates and terminates the tracks with some delay, while the smoother does this instantaneously, and that the state estimates from the smoother are slightly improved over that of the filter. These observations are confirmed by the results of 1000 Monte Carlo trials. The values of the OSPA distance for $c = 100m$ are shown in Figures 3 and 4 showing the total OSPA, along with its localization and cardinality components. The performance of the 1 and 2 lag smoothers is also shown. The miss distance spikes when the target appears and disappears, and maintains a value consistent with the measurement noise in between these times. The smoothers can be seen to perform better with increasing lag, and all perform better than the filter, incurring a lower error when the target appears and disappears, and also for the duration of the target existence. Over the entire scenario, the filter averages an error of $27.9m$, the 1-lag smoother $23.6m$, the 2-lag smoother $21.4m$, and the 3-lag smoother $20.6m$.

V. Discussion

In this paper, we demonstrate the first practical implementation of Bayesian smoothing using Finite Set Statistics, in the context of the joint detection and estimation of a single target. In particular, the forward backward smoother was derived using Finite Set Statistics, and a particle implementation was

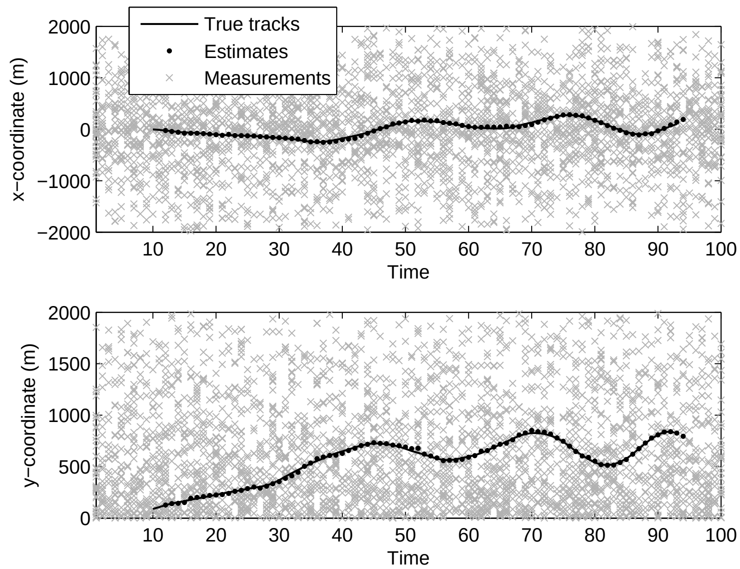

Fig. 1. Standard Filter Results

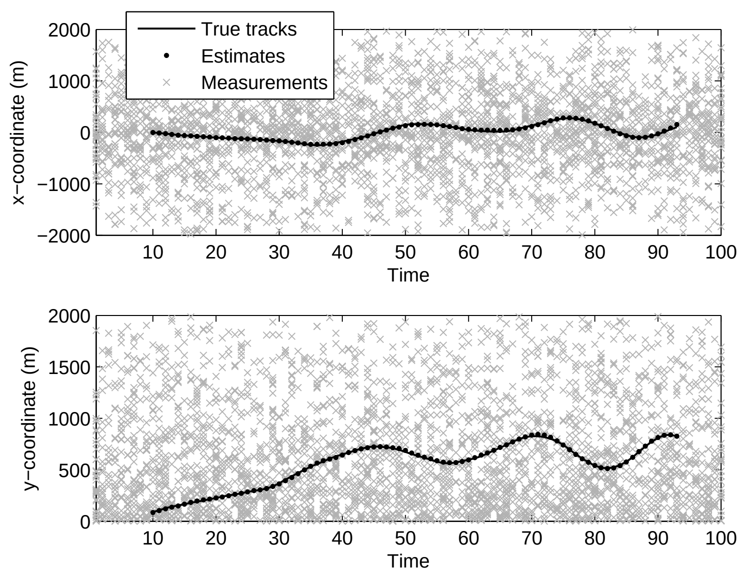

Fig. 2. Forward-Backward 3-Lag Smoother Results

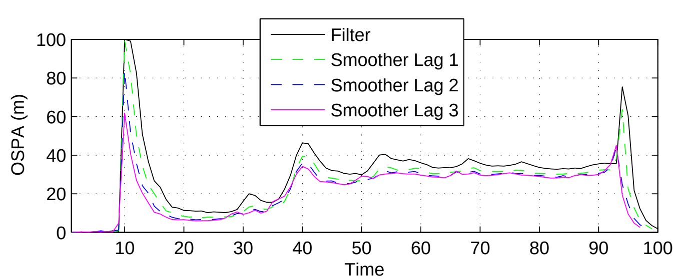

Fig. 3. OSPA Distance for filter and smoothers

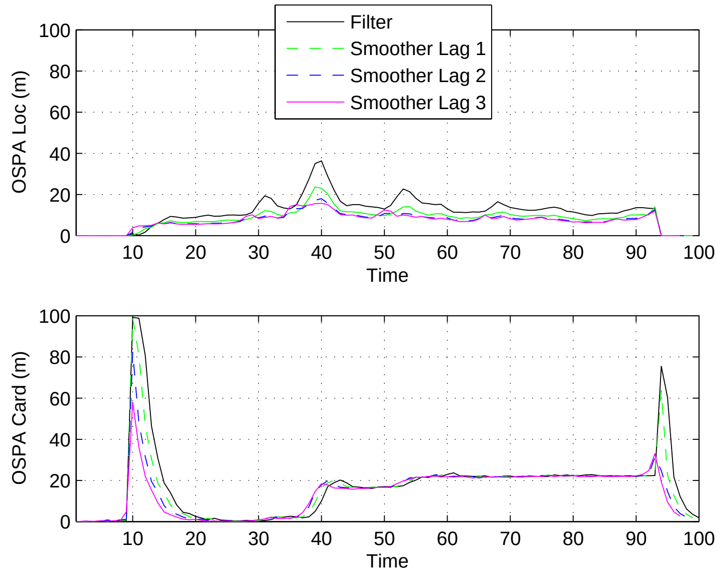

Fig. 4. OSPA Localization and Cardinality Components for filter and smoothers

given. We have also demonstrated that the use of smoothing can lead to an improvement in performance, both in the identification of whether there is a target present, and if so, in the state estimate itself.

References

[1] R. E. Kalman. A new approach to linear filtering and prediction problems. *Transactions of the ASME–Journal of Basic Engineering*, 82(Series D):35–45, 1960.

[2] B. D. Anderson and J. B. Moore. *Optimal Filtering*. Prentice-Hall, New Jersey, 1979.

[3] J. Meditch. A survey of data smoothing for linear and nonlinear dynamic systems. *Automatica*, 9:151162, 1973.

[4] A. Doucet, N. de Freitas, and N. Gordon. Sequential Monte Carlo Methods in Practice. *Springer-Verlag*, 2001.

[5] D.E. Clark, B.-T. Vo, and B.-N. Vo. Forward-backward sequential monte carlo smoothing for joint target detection and tracking. In *Information Fusion, 2009. FUSION '09. 12th International Conference on*, pages 899 –906, 2009.

[6] R. Mahler. Multi-target Bayes filtering via first-order multi-target moments. *IEEE Trans. AES*, 39(4):1152–1178, 2003.

[7] R. P. S. Mahler. Statistical Multisource Multitarget Information Fusion. *Artech House*, 2007.

[8] B.-T. Vo, B.-N. Vo, and A. Cantoni. Bayesian filtering with random finite set observations. *IEEE Transactions on Signal Processing*, 56(4):1313–1326, 2008.

[9] B.-T. Vo. *Random Finite Sets in Multi-Object Filtering*. PhD thesis, University of Western Australia, 2008.

[10] G. Kitagawa. Non-Gaussian state-space modeling of nonstationarytime series. *Journal American Statistical Association*, 82:1032–1063, 1987.

[11] B.-N. Vo, S. Singh, and A. Doucet. Sequential Monte Carlo methods for Multi-target Filtering with Random Finite Sets. *IEEE Trans. Aerospace Elec. Systems*, 41, No.4:1224–1245, 2005.

[12] D. Schuhmacher, B.-T. Vo, and B.-N. Vo. A consistent metric for performance evaluation of multi-object filters. *IEEE Transactions on Signal Processing*, 56(8):3447–3457, 2008.

Acknowledgements

Dr. Ba Tuong Vo is the recipient of an Australian Research Council Post Doctoral Fellowship under Discovery Project DP0989007. Dr. Daniel Clark is a Royal Academy of Engineering/EPSRC Research

Fellow supported by EP/H010866/1. Dr. Ba-Ngu Vo is supported under the Australian Research Council's Discovery Project funding scheme DP0880553.