Proceedings of the 45th IEEE Conference on Decision & Control
Manchester Grand Hyatt Hotel
San Diego, CA, USA, December 13-15, 2006

Convolution particle filtering for parameter estimation in general state-space models

Fabien Campillo & Vivien Rossi
INRIA/IRISA
Campus de Beaulieu
35042 Rennes cedex, France
Fabien.Campillo@inria.fr Vivien.Rossi@inria.fr

***Abstract*— The state-space modeling of partially observed dynamic systems generally requires estimates of unknown parameters. From a practical point of view, it is relevant in such filtering contexts to simultaneously estimate the unknown states and parameters.**

Efficient simulation-based methods using convolution particle filters are proposed. The regularization properties of these filters is well suited, given the context of parameter estimation. Firstly the usual non Bayesian statistical estimates are considered: the conditional least squares estimate (CLSE) and the maximum likelihood estimate (MLE). Secondly, in a Bayesian context, a Monte Carlo type method is presented. Finally we present a simulated case study.

***Index Terms*— Hidden Markov model, parameter estimation, particle filter, convolution kernels, conditional least squares estimate, maximum likelihood estimate**

I. INTRODUCTION

Consider a general state-space dynamical system described by an unobserved state process x_t and an observation process y_t taking values in $\mathbb{R}^d$ and $\mathbb{R}^q$ respectively. This system depends on an unknown parameter $\theta \in \mathbb{R}^p$. Suppose that the state process is Markovian, and that the observations y_t are independent conditionally to the state process. Suppose also that the distribution law of y_t depends only on x_t. Hence this system is completely described by the state process transition density and the emission density, namely

$$\begin{aligned} x_t|x_{t-1} &\sim f_t(x_t|x_{t-1}, \theta)\,, \\ y_t|x_t &\sim h_t(y_t|x_t, \theta)\,, \end{aligned} \tag{1}$$

and by the initial density law π_0 of x_0.

The goal is to estimate simultaneously the parameter θ and the state process x_t based on the observations $y_{1:t} = \{y_1, \ldots, y_t\}$.

In the nonlinear hidden processes framework, the parameter estimation procedure is often based on an approximation of the optimal filter. The extended Kalman filter and its various alternatives can give good results in practice but suffer from an absence of theoretical backing. The particle filters propose a good alternative: in many practical cases they give better results, moreover their theoretical properties are becoming increasingly well understood [1] [2] [3].

It is thus particularly appealing to use particle filtering in order to estimate parameters in partially observed systems. For a review of the question, one can consult [4] or [5]. There are two main approaches:

- The non Bayesian approach which consists of minimizing a given cost function like the conditional least squares criterion or by maximizing the likelihood function. These methods are usually performed in batch processes but can also be extended to recursive procedures.
- The Bayesian approach where an augmented state variable which includes the parameter is processed by a filtering procedure. These methods suppose that a prior law is given for the parameter and are performed online.

In practice, the first approach could be used as an initialization for the second one.

Due to the partially observed system framework, the objective function introduced in the first approach should be approximated for various values of the parameter θ. This is done via the particle approximation of the conditional law $p(y_t|y_{1:t-1}, \theta)$. The Monte Carlo nature of this particle approximation will make the optimization problematic. However, recent analyses propose significant improvements of these aspects [6] [4].

The second approach takes place in a classical Bayesian framework, a prior probability law $\rho(\theta)$ is thus introduced on the parameter θ. A new state variable (x_t, θ_t), joining all the unknown quantities, is considered and the posterior law $p(x_t, \theta_t|y_{1:t})$ is then approximated using particle filters.

In this paper we propose and compare different estimates corresponding to these two approaches and based on convolution particle filter introduced in [7].

II. THE CONVOLUTION FILTERS

To present the convolution filter, suppose that the parameter θ is known and consider:

$$\begin{aligned} x_t|x_{t-1} &\sim f_t(x_t|x_{t-1})\,, \\ y_t|x_t &\sim h_t(y_t|x_t)\,. \end{aligned} \tag{2}$$

The objective is to estimate recursively the optimal filter

$$p(x_t|y_{1:t}) = \frac{p(x_t, y_{1:t})}{p(y_{1:t})} = \frac{p(x_t, y_{1:t})}{\int p(x_t, y_{1:t})\, dx_t} \tag{3}$$

where $p(x_t, y_{1:t})$ is the $(x_t, y_{1:t})$ joint density.

1-4244-0171-2/06/$20.00 ©2006 IEEE.

Assumption: Suppose that we know how to sample from the laws $f_t(\cdot|x_{t-1})$, $h_t(\cdot|x_t)$ and also from the initial law π_0.

Note that the explicit description of the conditional densities f_t and h_t is useless whereas for the standard particle filtering approaches h_t should be stated explicitly. For example, in case of observation equations like $y_t = H(x_t, v_t)$ or $H(x_t, y_t, v_t) = 0$, where v_t is a noise, the conditional density h_t is in general not available.

A. The simple convolution filter (CF)

Let $\{x_0^i\}_{i=1\cdots n}$ be a sample of size n of π_0. For all $i = 1\cdots n$, starting from x_0^i, t successive simulations from the system (2) lead to a sample $\{x_t^i, y_{1:t}^i\}_{i=1\cdots n}$ from $p(x_t, y_{1:t})$. We get the following empirical estimate of the joint density:

$$p(x_t, y_{1:t}) \simeq \frac{1}{n}\sum_{i=1}^{n} \delta_{(x_t^i, y_{1:t}^i)}(x_t, y_{1:t}) \tag{4}$$

where δ_x is the Dirac measure in x.

The Kernel estimate $p_t^n(x_t, y_{1:t})$ of $p(x_t, y_{1:t})$ is then obtained by convolution of the empirical measure (4) with an appropriate kernel (cf. Appendix I):

$$p_t^n(x_t, y_{1:t}) \stackrel{\text{def}}{=} \frac{1}{n}\sum_{i=1}^{n} K_{h_n}^x(x_t - x_t^i)\; K_{h_n}^{\bar y}(y_{1:t} - y_{1:t}^i)$$

where $K_{h_n}^{\bar y}(y_{1:t} - y_{1:t}^i) \stackrel{\text{def}}{=} \prod_{s=1}^t K_{h_n}^y(y_s - y_s^i)$ in which $K_{h_n}^x$, $K_{h_n}^y$ are Parzen-Rosenblatt kernels of appropriate dimensions. Note that in $K_{h_n}^x(x_t - x_t^i)$ (resp. $K_{h_n}^y(y_t - y_t^i)$) h_n could implicitly depend on n, d and $x_t^{1:n}$ (resp. n, q and $y_t^{1:n}$) (see Section II-C).

From (3), an estimate of the optimal filter is then:

$$p_t^n(x_t|y_{1:t}) \stackrel{\text{def}}{=} \frac{\sum_{i=1}^n K_{h_n}^x(x_t - x_t^i)\; K_{h_n}^{\bar y}(y_{1:t} - y_{1:t}^i)}{\sum_{i=1}^n K_{h_n}^{\bar y}(y_{1:t} - y_{1:t}^i)} \tag{5}$$

The basic convolution filter (CF) is defined by the density estimate (5). A simple recursive algorithm for its practical computation is presented in Table I.

Convergence properties of $p_t^n(x_t|y_{1:t})$ to the optimal filter are ensured [7] when $h_n \to 0$ and $n h_n^{tq+d} \to \infty$. Just like the Monte Carlo filters without resampling, it implies that n must grow with t to maintain a good estimation. A better approach with a resampling step is proposed in the next section.

B. The resampled convolution filter (R-CF)

A resampling step can take place very easily at the beginning of each time step of the basic CF algorithm, the resulting procedure is presented in Table II

C. Comments

The practical use of the CF and R-CF filters requires the choice of the kernel functions K^x, K^y and of the bandwidth parameters h_n^x, h_n^y. The nature of the kernel does not appreciably affect the quality of the results.

The choice $h_n^x = C_x \times n^{-1/(4+d)}$, $h_n^y = C_y \times n^{-1/(4+q)}$ is optimal for the mean square error criterion. The choice of the C's is a critical issue for density estimation and sophisticated techniques have been proposed (see, e.g., [8]). In the on-line context of nonlinear filtering these techniques are not usable. Moreover, particle filtering is aimed to "track" the state and not really to sharply estimate the conditional density.

The generic form $C_x = c_x \times [\mathrm{Cov}(x_t^1, \ldots, x_t^n)]^{1/2}$, $C_y = c_y \times [\mathrm{Cov}(y_t^1, \ldots, y_t^n)]^{1/2}$ with c_x, $c_y \simeq 1$ gives good results. For the simulations of the last section, we take a Gaussian kernel and we will see that the c's are easily adjusted.

for $t = 0$
  initial sampling: $x_0^1 \cdots x_0^n \sim \pi_0$
  weight initialization: $w_0^i \leftarrow 1$ for $i = 1:n$
for $t \geq 1$
  for $i = 1:N$
    state sampling: $x_t^i \sim f_t(\cdot|x_{t-1}^i)$
    observation sampling: $y_t^i \sim h_t(\cdot|x_t^i)$
    weight updating: $w_t^i \leftarrow w_{t-1}^i\, K_{h_n}^y(y_t - y_t^i)$
  filter updating: $p_t^n(x_t|y_{1:t}) = \frac{\sum_{i=1}^n w_t^i\, K_{h_n}^x(x_t - x_t^i)}{\sum_{i=1}^n w_t^i}$

TABLE I
THE SIMPLE CONVOLUTION FILTER (CF).

for $t = 0$
  filter initialization: $p_0^n \leftarrow \pi_0$
for $t \geq 1$
  resampling: $\bar x_{t-1}^1 \cdots \bar x_{t-1}^n \sim p_{t-1}^n$
  state sampling: $x_t^i \sim f_t(\cdot|\bar x_{t-1}^i)$ for $i = 1:n$
  observation sampling: $y_t^i \sim h_t(\cdot|x_t^i)$ for $i = 1:n$
  filter updating:
  $$p_t^n(x_t|y_{1:t}) = \frac{\sum_{i=1}^n K_{h_n}^y(y_t - y_t^i)\; K_{h_n}^x(x_t - x_t^i)}{\sum_{i=1}^n K_{h_n}^y(y_t - y_t^i)}$$

TABLE II
THE RESAMPLED CONVOLUTION FILTER (R-CF).

III. CONDITIONAL LEAST SQUARES ESTIMATE

The standard least squares estimate is not obtainable here since only the y_t's are available and, moreover, they are dependent. Thus let us consider a conditional least squares estimate, introduced to treat the time series, see [9].

Let $\{y_t\}_{t\geq 1}$ the stochastic process defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P}_\theta)$, whose distribution depends on the parameter $\theta \in \mathbb{R}^p$. Let θ_* the true value of the parameter. The conditional least squares estimate of θ is the value $\hat\theta_T$ which minimizes

$$Q_T(\theta) \stackrel{\text{def}}{=} \sum_{t=1}^{T} |y_t - \hat y_t(\theta)|^2 \text{ with } \hat y_t(\theta) \stackrel{\text{def}}{=} \mathbb{E}_\theta[y_t|y_{1:t-1}] \tag{6}$$

where $\mathbb{E}_\theta[y_1|y_{1:0}] = \mathbb{E}_\theta[y_1]$. In general, and especially in our context, the quantity $\mathbb{E}_\theta[y_t|y_{1:t-1}]$ is unreachable. It can be estimated using a particle filter. The conditional density of

y_t given $y_{1:t-1}$ is

$$p(y_t|y_{1:t-1},\theta) = \frac{p(y_{1:t}|\theta)}{p(y_{1:t-1}|\theta)} = \frac{p(y_{1:t}|\theta)}{\int p(y_{1:t}|\theta)\,dy_t}$$

so that

$$\hat{y}_t(\theta) = \frac{\int y_t\, p(y_{1:t}|\theta)\,dy_t}{\int p(y_{1:t}|\theta)\,dy_t} \tag{7}$$

For θ and $t \geq 1$ given, it is possible to generate n trajectories $(x^i_{0:t}, y^i_{1:t})$, for $i = 1\cdots n$, according to (1). Finally $\frac{1}{n}\sum_{i=1}^n y^i_t\, K^{\bar{y}}_{h_n}(y_{1:t-1} - y^i_{1:t-1})$ and $\frac{1}{n}\sum_{i=1}^n K^{\bar{y}}_{h_n}(y_{1:t-1} - y^i_{1:t-1})$ are respectively the convolution kernel estimates of the numerator and denominator in (7). Hence the estimate of $\hat{y}_t(\theta)$ built from these n trajectories is

$$\hat{y}^n_t(\theta) \stackrel{\text{def}}{=} \frac{\sum_{i=1}^n y^i_{t+1}\, K^{\bar{y}}_{h_n}(y_{1:t-1} - y^i_{1:t-1})}{\sum_{i=1}^n K^{\bar{y}}_{h_n}(y_{1:t-1} - y^i_{1:t-1})}\,.$$

We take

$$\hat{Q}^n_T(\theta) \stackrel{\text{def}}{=} \sum_{t=1}^T |y_t - \hat{y}^n_t(\theta)|^2 \tag{8}$$

to estimate the function Q_T. The associated least squares estimate is then $\hat{\theta}^n_T = \arg\min_\theta \hat{Q}^n_T(\theta)$.

IV. Maximum likelihood estimate

The likelihood function is by definition:

$$L_T(\theta) \stackrel{\text{def}}{=} p(y_{1:T}|\theta) = p(y_1|\theta)\prod_{t=2}^T p(y_t|y_{1:t-1},\theta)\,. \tag{9}$$

Of course this function is not generally computable, it is then necessary to have recourse to estimation, see [10] [11]. The practical likelihood estimate depends on the type of convolution filter used.

A. Maximum likelihood estimation with the CF

In the CF case an immediate estimate is:

$$\hat{L}^n_T(\theta) \stackrel{\text{def}}{=} p^n(y_{1:T}|\theta) = \frac{1}{n}\sum_{i=1}^n K^{\bar{y}}_{h_n}(y_{1:T} - y^i_{1:T})$$

Thus $\hat{\theta}^n_T = \arg\max_\theta \hat{L}^n_T(\theta)$ approximates the maximum likelihood estimate.

B. Maximum likelihood estimation with the R-CF

For the R-CF formalization is not immediate. However all the quantities necessary to compute an estimate are made available with the R-CF algorithm. Indeed, the variables $\{y^i_{t+1}\}_{i=1\cdots n}$ generated in the observation sampling step of the R-CF algorithm are realizations of $p^n(y_{t+1}|y_{1:t},\theta)$. Thus by applying a convolution kernel to $\{y^i_{t+1}\}_{i=1\cdots n}$, we obtain the following estimate of the likelihood function:

$$\hat{L}^{n,r}_T(\theta) = \prod_{t=1}^T \frac{1}{n}\sum_{i=1}^n K^y_{h_n}(y_t - y^i_t) \tag{10}$$

and $\hat{\theta}^{n,r}_T = \arg\max_\theta \hat{L}^{n,r}_T(\theta)$ approximates the maximum likelihood estimate.

V. Optimization difficulties

For each fixed value of θ, approximations (8) and (10) of the least squares function (6) and of the likelihood function (9) are computed through kernel particle filters with sampling procedures based on laws depending on θ. These approximations will not be as smooth as their original counterparts and standard optimization procedures will severely fail in such a context. Therefore, it is necessary to use specific optimization techniques.

This issue can be addressed by stochastic approximation and optimization methods. Recently Doucet [4] proposed a Robbins-Monroe procedure in this HMM framework. The principal defect of these approaches is the slowness of their convergence, in spite of the efforts to improve this aspect, the computing times remain high in practice.

When the random quantities in the dynamic system, generally the noises, are independent of the other quantities, it is possible to freeze their values to one of their realizations so that the functions to optimize in θ is not stochastic any more. This technique can only be applied to CF filter, indeed for the R-CF it is impossible to freeze the resampling steps. Hence, because of the particle impoverishment of the CF filter described above, this algorithm is only valid for short length time series.

This approach is connected with techniques of optimization on MCMC estimates. The principle is as follows: for every time t, the simulated random quantities are frozen to their realizations, it is then possible to use the traditional minimization algorithms like Gauss-Newton. The parameter estimation is thus obtained for a given random realization. The study of this method for static optimization is carried out in [12].

An adaption to the sequential context of nonlinear filtering, for the maximization of the likelihood, is proposed in [13]. Several problems arise in practice, for example, for some values of the parameters, all the particle weights can be low providing a poor quality estimate. This approach remains extremely attractive as it then becomes possible to carry out optimizations using only one sample and consequently is valid for small variations of the parameter value. Thus Cérou et al [6] proposed an estimate of the derived filter based on this principle.

Of course, it is also possible to use a stochastic version of EM algorithm for this type of problem of optimization. Some references for this alternative are proposed in [13], but the difficulty of implementation makes it unfeasible.

VI. R-CF with unknown parameters approach

Suppose that the parameter θ is a random variable with a given prior law $\rho(\theta)$ and consider the augmented state variable (x_t, θ_t) with the following dynamic:

$$\theta_t = \theta_{t-1}\,, \qquad \theta_0 \sim \rho\,, \tag{11a}$$

$$x_t|x_{t-1} \sim f_t(x_t|x_{t-1},\theta_t)\,, \qquad x_0 \sim p_0\,, \tag{11b}$$

$$y_t|x_t \sim h_t(y_t|x_t,\theta_t)\,. \tag{11c}$$

The posterior law of θ_t is then given by the nonlinear filter.

The constant dynamic (11a) may lead to the divergence of the standard particle filters. This is due to the fact that the parameter space is only explored at the initialization step of the particle filter which causes the impoverishment of the variety of the relevant particles. Among the approaches proposed to avoid this trouble, Storvik [14] marginalizes parameters out of the posterior distribution then assume that the concerned parameters depend on sufficient statistics which allows their simulations and avoids the impoverishment. However it is not practically useful for general systems. Kitagawa [10] and Higuchi [15] set an artificial dynamic on the parameter, like $\theta_t = \theta_{t-1} + \zeta_t$ or more complex, thus risking mismatching the system dynamic. Gilks & Berzuini [16], Lee & Chia [17] add a Markov chain Monte Carlo procedure to increase the particle diversity, but this is cumbersome. To avoid these additions West [18], Liu & West [5] propose to smooth the empirical measure of the parameter posterior law with a Gaussian distribution.

More generally, regularization techniques are used to avoid the degeneration of the particle filters. Most of the time the regularization only concerns the state variables, see [19] and [20]. However this approach still suffers from some of the restrictions of the traditional methods: it requires the non nullity of the noise variances and the analytical availability of the likelihood function. These restrictions were dropped in [21] by the regularization of the observation model. However, as the state model is not regularized, the approach remains sensitive to the problem of degeneration of the particle filters.

In order to circumvent these problems, Rossi & Vila [7] jointly regularized the state model and the observation model. Their approach can be interpreted as a generalization of the preceding models, thanks to an extension of the concept of particle which includes the state and the observation. However, the construction and the theoretical study of the corresponding filters are different as they are based on the nonparametric estimate of the conditional densities by convolution kernels. The filter used in this section to estimate simultaneously the state and the parameters in (11), extends the results of [7]. It is not necessary for the kernel to be Gaussian as in West [18], any kernel satisfying the conditions of the Appendix I will be valid.

The regularization with convolution kernels can also be viewed as artificial noise. Thus our approach is connected to the methods [10], [15] presented previously. However contrary to these methods, it respects dynamics (11a) and allows convergence results. In terms of artificial noise on dynamics, this means that we have identified a whole family of acceptable noises and that we have also characterized the way in which their variance must decrease to zero.

The R-CF filter (Table II) applied to the system (11) leads to the algorithm presented in Table III. It provides consistent estimates of $p(x_t, \theta_t|y_{1:t})$, $p(x_t|y_{1:t})$ and $p(\theta_t|y_{1:t})$. The first probability law is the key element of the algorithm. It is used as a sample generator and it is updated at every time iteration. The two last laws are used to estimate the state x_t and the parameter θ_t respectively.

In practice, the parameter prior law $\rho(\theta)$, the number of particles n, the kernels K and the associated bandwidth parameters h_n must by chosen by the user.

Generate $\bar{x}_0^i \sim p(x_0)$ and $\bar{\theta}_0^i \sim \rho(\theta)$ for $i = 1 \cdots n$

For $t = 1$

generation of the trajectories: for $i = 1 \cdots n$

$x_1^i \sim f_1(\cdot|\bar{x}_0^i, \bar{\theta}_0^i)$
$y_1^i \sim h_1(\cdot|x_1^i, \bar{\theta}_0^i)$
$\theta_1^i = \bar{\theta}_0^i$

estimate of the densities:

$$p_1^n(x_1, \theta_1|y_1) = \frac{\sum_{i=1}^n K_{h_n}^y(y_1 - y_1^i)\, K_{h_n}^\theta(\theta_1 - \theta_1^i)\, K_{h_n}^x(x_1 - x_1^i)}{\sum_{i=1}^n K_{h_n}^y(y_1 - y_1^i)}$$

$$p_1^n(\theta_1|y_1) = \frac{\sum_{i=1}^n K_{h_n}^y(y_1 - y_1^i)\, K_{h_n}^\theta(\theta_1 - \theta_1^i)}{\sum_{i=1}^n K_{h_n}^y(y_1 - y_1^i)}$$

$$p_1^n(x_1|y_1) = \frac{\sum_{i=1}^n K_{h_n}^y(y_1 - y_1^i)\, K_{h_n}^x(x_1 - x_1^i)}{\sum_{i=1}^n K_{h_n}^y(y_1 - y_1^i)}$$

For $t \geq 2$

generation of the trajectories: for $i = 1 \cdots n$

$(\bar{x}_{t-1}^i, \bar{\theta}_{t-1}^i) \sim p_{t-1}^n(x_{t-1}, \theta_{t-1}|y_{1:t-1})$
$x_t^i \sim f_t(\cdot|\bar{x}_{t-1}^i, \bar{\theta}_{t-1}^i)$
$y_t^i \sim h_t(\cdot|x_t^i, \bar{\theta}_{t-1}^i)$
$\theta_t^i = \bar{\theta}_{t-1}^i$

estimate of the densities:

$$p_t^n(x_t, \theta_t|y_{1:t}) = \frac{\sum_{i=1}^n K_{h_n}^y(y_t - y_t^i)\, K_{h_n}^\theta(\theta_t - \theta_t^i)\, K_{h_n}^x(x_t - x_t^i)}{\sum_{i=1}^n K_{h_n}^y(y_t - y_t^i)}$$

$$p_t^n(\theta_t|y_{1:t}) = \frac{\sum_{i=1}^n K_{h_n}^y(y_t - y_t^i)\, K_{h_n}^\theta(\theta_t - \theta_t^i)}{\sum_{i=1}^n K_{h_n}^y(y_t - y_t^i)}$$

$$p_t^n(x_t|y_{1:t}) = \frac{\sum_{i=1}^n K_{h_n}^y(y_t - y_t^i)\, K_{h_n}^x(x_t - x_t^i)}{\sum_{i=1}^n K_{h_n}^y(y_t - y_t^i)}$$

TABLE III

THE RESAMPLED CONVOLUTION FILTER FOR BAYESIAN ESTIMATION.

VII. SIMULATED CASE STUDIES: BEARINGS–ONLY TRACKING

We compare the convolution filter (R-CF) with the standard bootstrap particle filter (BPF) and with the extended Kalman filter (EKF) applied to the classical problem of bearings-only tracking in the plane [22].

Consider a mobile (the target) with a rectilinear uniform motion (i.e., with constant bearing and speed) in the plane. This mobile is tracked by an observer with a given trajectory. The state vector $x_t = (p_t^1, p_t^2, v_t^1, v_t^2)^*$ represents the relative positions and velocities vector of the Cartesian coordinates for the difference between the tracked object and the observer $x_t = x_t^{\text{tg}} - x_t^{\text{obs}}$. This state vector is solution of a linear noise–free system:

$$x_{t+1} = A_t\, x_t + B_t \tag{12}$$

where A_t and B_t are given. The observations are a sequence of bearings corrupted by noise:

$$y_t = \tan^{-1}(p_t^1/p_t^2) + \sigma\, v_t$$

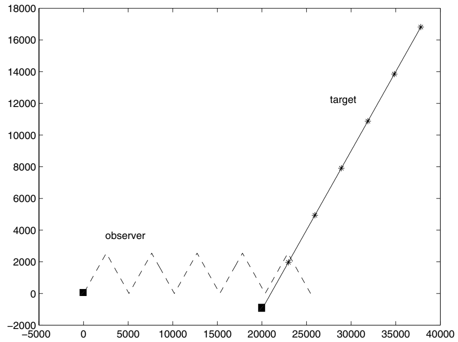

Fig. 1. Simulation scenario. Total observation time 1 hour, sampling interval 4s. Initial relative distance 20025m, target speed 7m/s, observer speed 10m/s. Trajectories: target (plain line), maneuvering observer (dashed line), initial positions (black squares).

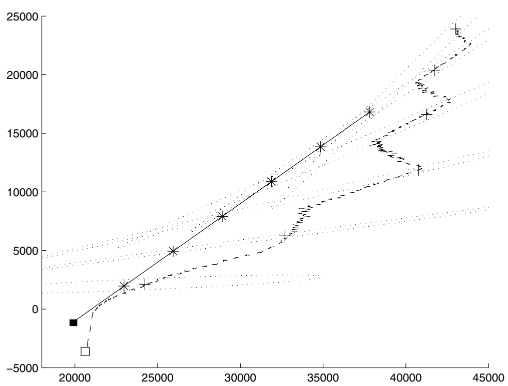

Fig. 2. Extended Kalman filter. Plain line: true trajectory. Dashed line: empirical estimated trajectory after 15 Monte Carlo independent runs and the corresponding empirical uncertainty ellipses (every 10 minutes).

where v_t is a white Gaussian noise $\mathcal{N}(0,1)$ and $\sigma = 0.5$ degree. The simulation scenario is described in Fig. 1. The initial state law is $(p_0^1, p_0^2, v_t^1, v_t^2)^* \sim \mathcal{N}\big((23000, -3000, -10, 0)^*, \mathrm{diag}(5000^2, 5000^2, 10^2, 10^2)\big)$ while the true value is $(20000, -1000, -12, -2)^*$.

We perform 15 independent Monte Carlo runs of this scenario. In Figs. 2 to 6 we present the corresponding empirical position (the empirical estimated trajectory) and the corresponding empirical uncertainty ellipses (every 10 minutes). For R-CF and BPF we use 10000 particles. Calculation times and memory requirements are equivalent for R-CF and BPF.

This example is known to be unfavorable for EKF but it does show the advantage of our approach. Moreover, the standard particle filter requires the addition of an artificial noise in the state equation (12). The adjustment of the intensity of this noise is complicated, so it is a delicate process implementing the standard particle filter, see Figs. 3 and 4. Filter R-FC appears simpler and more robust in all the cases.

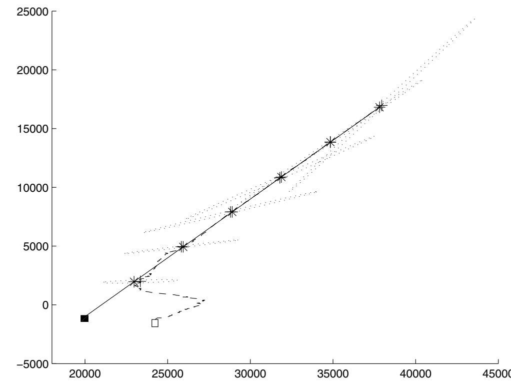

Fig. 3. Bootstrap particle filter (BPF) with artificial Gaussian noise on (12): 0.025m/s standard deviation on the velocity components and 25m on the position components.

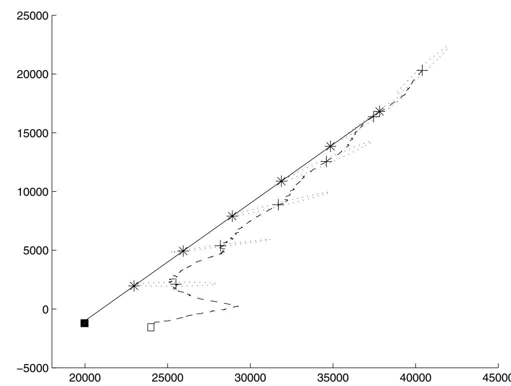

Fig. 4. Bootstrap particle filter (BPF) with artificial Gaussian noise on (12): 0.05m/s standard deviation on the velocity components and 50m on the position components.

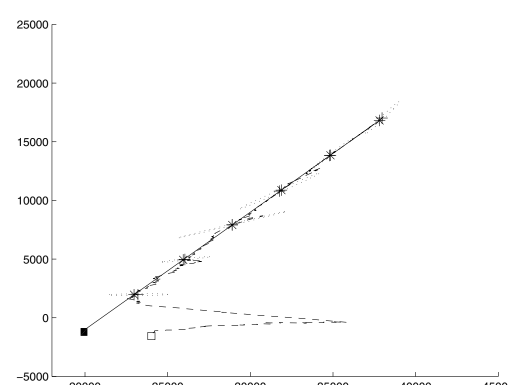

Fig. 5. Resampled convolution filter (RCF): $c_x = 0.8$ and $c_y = 1$.

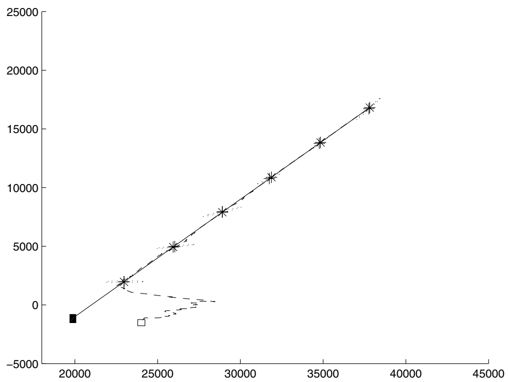

Fig. 6. Resampled convolution filter (RCF): $c_x = 0.6$ and $c_y = 1$.

VIII. Conclusion and Discussion

The different estimation approaches proposed show the large potential of the convolution filters.

The first approach, based on the maximization of the likelihood estimate and the minimization of the conditional least squares estimate, presents several drawbacks. From the practical point of view they require a high computation time and the choice of the number of observation T. From the theoretical point of view, their convergence is ensured under uniform convergence assumptions, which is difficult to verify for a given dynamic system. However, the convolution filters approach is a good alternative to the stochastic optimization, and can be used to perform the initialization of a Bayesian procedure.

The R-CF with unknown parameters approach introduced in the last section is perfectly suited for online estimation and their theoretical properties are clearly established without need of additional strong assumptions. Thus this last approach is interesting especially if the primary objective is the filtering in spite of uncertainties with the model.

Appendix I
Kernel Estimation

A kernel $K : \mathbb{R}^d \mapsto \mathbb{R}$ is a bounded, positive, symmetric application such that $\int K(x)\,dx = 1$. We denote

$$K_{h_n}(x) \stackrel{\text{def}}{=} \tfrac{1}{h_n^d}\,K\big(\tfrac{x}{h_n}\big)\,.$$

$h_n > 0$ is the bandwidth parameter. The Gaussian kernel is $K(x) = \big(\frac{1}{\sqrt{2\pi}}\big)^d\, e^{-|x|^2/2}$. A Parzen-Rosenblatt kernel is a kernel such that $\|x\|^d K(x) \to 0$ as $\|x\| \to \infty$.

Let $X_1 \cdots X_n$ be i.i.d. random variables with common density f. The kernel estimator of f associated with the kernel K is given by

$$f_n(x) = \tfrac{1}{n\,h_n^d}\sum_{i=1}^n K\big(\tfrac{x-X_i}{h_n}\big) = (K_{h_n} * \mu_n)(x)$$

for $x \in \mathbb{R}^d$; $\mu_n = \frac{1}{n}\sum_{i=1}^n \delta_{X_i}$ is the empirical measure associated with $X_1 \cdots X_n$.

References

[1] P. Del Moral, *Feynman-Kac formulae – Genealogical and interacting particle approximations*. New York: Springer–Verlag, 2004.
[2] A. Doucet, N. de Freitas, and N. J. Gordon, Eds., *Sequential Monte Carlo Methods in Practice*. New York: Springer–Verlag, 2001.
[3] D. Crişan and A. Doucet, "A survey of convergence results on particle filtering for practitioners," *IEEE Transactions on Signal Processing*, vol. 50, pp. 736–746, 2002.
[4] A. Doucet and V. B. Tadić, "Parameter estimation in general state–space models using particle methods," *Annals of the Institute of Statistical Mathematics*, vol. 55, pp. 409–422, 2003.
[5] J. Liu and M. West, "Combined parameter and state estimation in simulation-based filtering," in *[2]*, ch. 12, pp. 197–223.
[6] F. Cérou, F. L. Gland, and N. J. Newton, "Stochastic particle methods for linear tangent filtering equations," in *Optimal Control and Partial Differential Equations*, J.-L. Menaldi, E. Rofman, and A. Sulem, Eds. Amsterdam: IOS Press, 2001, pp. 231–240.
[7] V. Rossi and J.-P. Vila, "Nonlinear filtering in discret time: A particle convolution approach," Biostatistic group of Montpellier, Technical Report 04-03, 2004.
[8] L. Devroye and G. Lugosi, *Combinatorial Methods in Density Estimation*. Springer–Verlag, 2001.
[9] H. Tong, *Nonlinear Time Series: A Dynamical System Approach*. Oxford: Oxford University Press, 1990.
[10] G. Kitagawa, "Non–Gaussian state space modeling of nonstationary time series (with discussion)," *Journal of the American Statistical Association*, vol. 82, pp. 1032–1063, 1987.
[11] ——, "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.
[12] C. J. Geyer, "Estimation and optimization of functions," in *Markov Chain Monte Carlo in practice*, D. S. W.R. Gilks, S. Richardson, Ed. London: Chapman & Hall, 1996, pp. 241–258.
[13] M. Hürzeler and H. Künsch, "Approximating and maximising the likelihood for a general state-space model," in *[2]*, ch. 12, pp. 159–175.
[14] G. Storvik, "Particle filters in state space models with the presence of unknown static parameters," *IEEE Transactions on Signal Processing*, vol. 50, pp. 281–289, 2002.
[15] T. Higuchi, "Monte Carlo filter using the genetic algorithm operators," *Journal of Statistical Computation and Simulation*, vol. 59, pp. 1–23, 1997.
[16] W. R. Gilks and C. Berzuini, "Following a moving target — Monte Carlo inference for dynamic Bayesian models," *Journal of the Royal Statistical Society, Series B*, vol. 63, pp. 127–146, 2001.
[17] D. S. Lee and N. K. K. Chia, "A particle algorithm for sequential Bayesian parameter estimation and model selection," *IEEE Transactions on Signal Processing*, vol. 50, pp. 326–336, 2002.
[18] M. West, "Approximating posterior distributions by mixtures," *Journal of the Royal Statistical Society, Series B*, vol. 55, pp. 409–422, 1993.
[19] P. Del Moral and L. Miclo, "Branching and interacting particle systems approximations of Feynman–Kac formulae with applications to nonlinear filtering," in *Séminaire de Probabilités XXXIV*, ser. Lecture Notes in Mathematics, J. Azéma, M. Émery, M. Ledoux, and M. Yor, Eds. Berlin: Springer–Verlag, 2000, vol. 1729, pp. 1–145.
[20] F. Le Gland and N. Oudjane, "A robustification approach to stability and to uniform particle approximation of nonlinear filters: the example of pseudo-mixing signals," *Stochastic Processes and their Applications*, vol. 106, pp. 279–316, 2003.
[21] P. Del Moral, J. Jacod, and P. Protter, "The Monte Carlo method for filtering with discrete–time observations," *Probability Theory and Related Fields*, vol. 120, pp. 346–368, 2001.
[22] B. Ristic, M. S. Arulampalam, and N. J. Gordon, *Beyond the Kalman Filter: Particle Filters for Tracking Applications*. Artech House, 2004.