PARTICLE FILTER AS A CONTROLLED MARKOV CHAIN FOR ON-LINE PARAMETER ESTIMATION IN GENERAL STATE SPACE MODELS

George Poyiadjis[1] *- Sumeetpal S. Singh*[1] *- Arnaud Doucet*[2]

[1] Dept of Engineering, University of Cambridge, Cambridge CB2 1PZ, UK.
[2]Depts of CS and Statistics, University of British Columbia, Vancouver, BC, Canada.
Email: gp243,sss40@cam.ac.uk - arnaud@stat.ubc.ca

ABSTRACT

In this paper we present a novel optimization method for on-line maximum likelihood estimation (MLE) of the static parameters of a general state space model. Our approach is based on viewing the particle filter as a controlled Markov chain, where the control is the unknown static parameters to be identified. The algorithm relies on the computation of the gradient of the particle filter using a score function approach.

1. INTRODUCTION

Despite the advent of Sequential Monte Carlo (SMC) filtering methods (aka particle filters), SMC based parameter estimation remains a challenging problem. The majority of the methods found in the literature augment the state to include the unknown static parameter and cast the problem as a filtering one [5], [10], [13]. These methods have their drawbacks [1]. Our approach is to solve the parameter estimation problem as a Recursive Maximum Likelihood (RML) problem. RML is well known but hitherto has been limited to finite state space models [8]. Recently, extensions to the general state space case have been proposed [3], [12], [1]. In [3] and [12], the authors solve the problem by directly approximating the derivative of the optimal filter, which is needed for their RML implementation, using SMC. In [1], novel on-line Expectation Maximization algorithms have been presented. The approach adopted in this paper is to view the particle filter as a controlled Markov chain. We define a suitable "average reward" performance criterion and optimize it. The control parameters of the Markov chain correspond to the unknown static parameters of the hidden Markov model to be identified. The approach we propose is an on-line method and we demonstrate good performance in several examples.

1.1. The State-Space Model

Let the hidden state process $\{X_n\}_{n\geq 0}$ and the observation process $\{Y_n\}_{n\geq 0}$ be $\mathbb{R}^{n_x}$ and $\mathbb{R}^{n_y}$ -valued stochastic processes defined on a measurable space $(\Omega, \mathcal{F})$. Let $\theta \in \Theta$ be the parameter vector where Θ is an open subset of $\mathbb{R}^m$. A general discrete-time state-space model represents the unobserved state $\{X_n\}_{n\geq 0}$ as a Markov process of initial density $X_0 \sim \mu$ and transition density $f_\theta(x_n|x_{n-1})$. The process $\{X_n\}_{n\geq 0}$ is unknown but is partially observed through the observation process $\{Y_n\}_{n\geq 0}$. The observations $\{Y_n\}_{n\geq 0}$ are assumed conditionally independent given $\{X_n\}_{n\geq 0}$ and admit a marginal density $g_\theta(y_n|x_n)$.

Consider for the time being a fixed parameter θ. We are interested in the so-called optimal filter that is defined as the posterior density of X_n given the observation set $Y_{0:n}$, denoted by $p_\theta(x_n|Y_{0:n})$[1].

[1]For any sequence $\{z_k\}$ and random process $\{Z_k\}$ we will use the notation $z_{i:j} = (z_i, z_{i+1}, ..., z_j)$ and $Z_{i:j} = (Z_i, Z_{i+1}, ..., Z_j)$.

Introducing a proposal distribution $q_\theta(x_n|Y_n, x_n - 1)$, which is needed for the SMC implementation below, the filtering density satisfies the recursion [4]

$$p_\theta(x_n|Y_{0:n}) \propto \int \alpha_\theta(x_{n-1:n}, Y_n)\, q_\theta(x_n|Y_n, x_{n-1})\, p_\theta(x_{n-1}|Y_{0:n-1})\, dx_{n-1} \quad (1)$$

$$\text{where } \alpha_\theta(x_{n-1:n}, Y_n) = \frac{g_\theta(Y_n|x_n)\, f_\theta(x_n|x_{n-1})}{q_\theta(x_n|Y_n, x_{n-1})}. \quad (2)$$

Note that the support of $q_\theta(x_n|Y_n, x_n - 1)$ should include that of $g_\theta(Y_n|x_n)\, f_\theta(x_n|x_n - 1)$. In general, (1) has no closed-form expression and particle methods are typically used to approximate it.

1.2. The SMC framework

The SMC method we will consider here approximates the optimal filtering density $p_\theta(x_{n-1}|Y_{0:n-1})$ by an set of $N \gg 1$ particles, i.e. $X_{n-1}^{(1:N)} \triangleq \left[X_{n-1}^{(1)}, \ldots, X_{n-1}^{(N)}\right]$ having equal weights. The filtering distribution at the next time step can be recursively approximated by a new set of particles $X_n^{(1:N)}$ generated via an importance sampling and a resampling step.

In the importance sampling step, a set of prediction particles $\widetilde{X}_n^{(1:N)}$ are generated independently using the proposal distribution introduced in (1),

$$\widetilde{X}_n^{(1:N)} \sim Q_\theta\left(\widetilde{x}_n^{(1:N)} \mid Y_n, X_{n-1}^{(1:N)}\right) = \prod_{i=1}^{N} q_\theta\left(\widetilde{x}_n^{(i)} | X_{n-1}^{(i)}, Y_n\right)$$

The i^{th} prediction particle is weighted by an importance weight $\widetilde{a}_{\theta,n}^{(i)}$ that corrects for the discrepancy with the "target" distribution. This is given by

$$a_{\theta,n}^{(i)} = \alpha_\theta\left(\widetilde{X}_n^{(i)}, X_{n-1}^{(i)}, Y_n\right), \quad \widetilde{a}_{\theta,n}^{(i)} = \frac{a_{\theta,n}^{(i)}}{\sum_{j=1}^{N} a_{\theta,n}^{(j)}}. \quad (3)$$

Note that the weight vector is a deterministic function of its arguments and will be denoted by

$$\widetilde{a}_{\theta,n}^{(1:N)} = \Phi_\theta\left(Y_{n-1}, \widetilde{X}_{n-1}^{(1:N)}, X_{n-2}^{(1:N)}\right). \quad (4)$$

In the resampling step, the particles $\widetilde{X}_n^{(1:N)}$ are multiplied or eliminated according to their importance weights $\widetilde{a}_{\theta,n}^{(1:N)}$ to give the new set of particles $X_n^{(1:N)}$, based on the mapping

$$X_n^{(1:N)} = H\left(\widetilde{X}_n^{(1:N)}, I_n^{(1:N)}\right) \triangleq [\underbrace{\widetilde{X}_n^{(1)}, \ldots, \widetilde{X}_n^{(1)}}_{I_n^{(1)} \text{ times}}, \ldots, \underbrace{\widetilde{X}_n^{(N)}, \ldots, \widetilde{X}_n^{(N)}}_{I_n^{(N)} \text{ times}}] \quad (5)$$

where $I_n^{(i)}$ represents the number of copies of particle $\widetilde{X}_n^{(i)}$. The resampling index vector $I_n^{(1:N)} \triangleq \left[I_n^{(1)}, \ldots, I_n^{(N)}\right] \sim L\left(\cdot \mid \widetilde{a}_{\theta,n-1}^{(1:N)}\right)$ is obtained using multinomial resampling, which is smooth and hence differentiable in the weights[2].

If we assume that at time $n-1$ a set of equally weighted particles $\widetilde{X}_{n-1}^{(1:N)} = \left[\widetilde{X}_{n-1}^{(1)}, \ldots, \widetilde{X}_{n-1}^{(N)}\right]$ is available, the full algorithm is summarized as follows:

Generic Sequential Monte Carlo algorithm (SIR)

Weighted resampling step

- Evaluate the importance weights $\widetilde{a}_{\theta,n-1}^{(i)}$ using (3).
- Sample the resampling index vector $I_{n-1}^{(1:N)} \sim L\left(\cdot \mid \widetilde{a}_{\theta,n-1}^{(1:N)}\right)$.
- Set $X_{n-1}^{(1:N)} = H\left(\widetilde{X}_{n-1}^{(1:N)}, I_{n-1}^{(1:N)}\right)$.

Importance sampling step

- For $i = 1, ..., N$, sample $\widetilde{X}_n^{(i)} \sim q_\theta\left(\cdot \mid Y_n, X_{n-1}^{(i)}\right)$.

2. PROBLEM FORMULATION

In a general state space model, the system evolves according to a true but unknown static parameter θ^*, i.e.

$$X_n | X_{n-1} = x_{n-1} \sim f_{\theta^*}(\,.\,|x_{n-1}) \tag{6}$$

$$Y_n | X_n = x_n \sim g_{\theta^*}(\,.\,|x_n). \tag{7}$$

The aim is to identify θ^* based on an infinite (or very large) observation sequence $\{Y_n\}_{n\geq 0}$, in an on-line fashion. Identification of θ^* can be achieved by maximizing, with respect to θ, the average log-likelihood function (see for example [9])

$$l(\theta) = \lim_{k\to\infty} \frac{1}{k+1} \sum_{n=0}^{k} \log p_\theta(Y_n \,|Y_{0:n-1}\,). \tag{8}$$

It can be shown that under suitable regularity conditions described in [14], this limit exists and $l(\theta)$ admits θ^* as a global maximum. The expression $p_\theta(Y_n \,|Y_{0:n-1}\,)$ can be written as

$$p_\theta(Y_n \,|Y_{0:n-1}\,) = \int\int \alpha_\theta\,(x_{n-1:n}, Y_n)\, q_\theta\,(x_n | \,Y_n, x_{n-1}) \times p_\theta\,(x_{n-1} | \,Y_{0:n-1})\, dx_{n-1} dx_n \tag{9}$$

and in general does not admit a closed form solution. However, it can be approximated numerically using the SMC algorithm of section 1.2. We propose here to maximize the alternative criterion that results from using a particle approximation for the likelihood in (8), i.e.

$$J(\theta) = \lim_{k\to\infty} \frac{1}{k+1} \sum_{n=0}^{k} \log \widehat{p}_\theta(Y_n \,|Y_{1:n-1}\,). \tag{10}$$

As detailed below, this approach amounts to treating the particle filter as a controlled Markov chain which is to be optimized.

2.1. Solution methodology

The general approach we will consider is based on the literature of controlled Markov chains [11]. Consider a Markov chain $\{Z_n\}_{n\geq 0}$, taking values in R^{n_z}, whose transition density $\mathbf{p}_\theta$ is parameterized by a tunable or control parameter $\theta \in \Theta$. We are interested in maximizing performance measures of the form $J(\theta) = \int \varphi_\theta(z)\lambda_\theta(z)dz$ w.r.t. θ, where φ_θ is some real valued function defined on R^{n_z} and λ_θ is the unique invariant distribution of the Markov chain. In most cases, $\lambda_\theta(z)$ is not known and $J(\theta)$ must be optimized based on knowledge of the transition density only. Ergodicity of $\{Z_n\}_{n\geq 0}$ is thus important. Specifically, the instantaneous expectation of the chain given by

$$J_n(\theta) \triangleq E_\theta\left\{\varphi_\theta(Z_n)\right\} = \int \varphi_\theta(z_n)\mathbf{p}_\theta(z_n\,|z_{n-1}\,)\ldots \times \mathbf{p}_\theta(z_1\,|z_0\,)\mathbf{p}\,(z_0)\, dz_{0:n} \tag{11}$$

converges under suitable ergodicity assumptions to the stationary performance measure, i.e. $\lim_{n\to\infty} J_n(\theta) = J(\theta)$. One way to exploit this in order to maximize $J(\theta)$ is to use a Stochastic Approximation (SA) algorithm that updates the parameter value at time n according to the recursion

$$\theta_n = \theta_{n-1} + \gamma_n \widehat{\nabla J}_n(\theta_{n-1}). \tag{12}$$

Here θ_{n-1} is the parameter estimate at time $n-1$ and $\widehat{\nabla J}_n$ denotes an estimate of ∇J_n, preferably unbiased. Under suitable conditions on the step size, θ_n will converge to $\vartheta^* = \arg\max_{\theta\in\Theta} J(\theta)$ (or to the set of maximizers in the case it is not unique).

In order to obtain unbiased estimates of $\nabla J_n(\theta)$, we will use the *score function* method, which is also known as the *likelihood ratio* [6]. The gradient of (11), applying the chain rule, is

$$\nabla J_n(\theta) = E\left\{(\nabla\varphi_\theta)\,(Z_n) + \varphi_\theta(Z_n)S_n(\theta)\right\},$$

where

$$S_n(\theta) = \sum_{m=1}^{n} \frac{\nabla \mathbf{p}_\theta(Z_m\,|Z_{m-1}\,)}{\mathbf{p}_\theta(Z_m\,|Z_{m-1}\,)} \tag{13}$$

is the score function. This leads to the following unbiased gradient estimate

$$\widehat{\nabla J}_n(\theta) = \nabla\varphi_\theta(Z_n) + \varphi_\theta(Z_n)S_n(\theta). \tag{14}$$

Note that the algorithm requires only a single realization or sample path of the Markov chain.

3. PARAMETER ESTIMATION OF SMC

We use the above ideas in the context of parameter estimation of SMC systems by defining an augmented system comprising of the hidden state, its observation and the particle filter, i.e.

$$Z_n = (X_n, Y_n, \widehat{p}_n),$$

where $\widehat{p}_n = \left[\widetilde{X}_n^{(1:N)}, X_{n-1}^{(1:N)}\right]$ is an approximation to the joint prediction density $q_\theta\,(x_n|\,Y_n, x_{n-1})\, p_\theta\,(x_{n-1}|\,Y_{1:n-1})$. A related approach was used in [2] to optimize the performance of SMC algorithms. It can be easily verified that Z_n satisfies the Markov property, since $\widehat{p}_n$ is a random function of $\widehat{p}_{n-1}$ and $Y_{n-1:n}$ only. Under some stability conditions, it can be further shown that Z_n is an ergodic Markov chain for each θ [14]. The performance criterion (10) is now given by[3]

$$J(\theta) \triangleq \iint_{\mathbb{R}^{n_y}\times\mathcal{P}(\mathbb{R}^{nx})} \varphi_\theta\,(y, \widehat{p}_n)\,\lambda_{\theta,\theta^*}\,(y, \widehat{p}_n)\; dy\, d\widehat{p}_n \tag{15}$$

where $\mathcal{P}(\mathbb{R}^{nx})$ is the space of probability distributions on $\mathbb{R}^{nx}$ and $\lambda_{\theta,\theta^*}(\cdot)$ is the invariant distribution of the observation process Y_n and the predictive particle filter $\widehat{p}_n$. Note that the invariant distribution $\lambda_{\theta,\theta^*}$ depends on the true parameter θ^* as well, because of the

[2] Differentiability is essential for the development of the algorithm. Resampling schemes that are not continuous w.r.t. the weight vector (e.g. residual and systematic) can still be used with appropriate smoothing.

[3] Note that the first component of Z_n is the unknown hidden state X_n of the true system. Therefore, we essentially consider cost functions $\varphi_\theta(\cdot)$ that do not depend on X_n.

hidden state and observation pair (X_n, Y_n). For maximum likelihood estimation we set the cost function to be

$$\varphi_\theta(Z_n) = \log\left(\frac{1}{N}\sum_{i=1}^{N}\alpha_\theta\left(X_{n-1}^{(i)}, \widetilde{X}_n^{(i)}, Y_n\right)\right). \quad (16)$$

Note that this is precisely the particle approximation $\log \widehat{p}_\theta(Y_n\,|Y_{1:n-1}\,)$ in (10).

Following the arguments in section 2.1, the solution of

$$\vartheta^* = \arg\max_{\theta\in\Theta} J\,(\theta) \quad (17)$$

can be obtained through the SA recursion in (12) that requires an unbiased estimate of ∇J_n. Before we calculate this, it is important to realize that in a general state space model, the evolution of X_n and Y_n are governed by the true parameter θ^* that we wish to identify, while the particle filter component of Z_n evolves according to a θ-dependent transition. Specifically, the transition density of the Markov chain $Z_n = (X_n, Y_n, \widehat{p}_n)$ is

$$X_n, Y_n, \widehat{p}_n | X_{n-1}, Y_{n-1}, \widehat{p}_{n-1} \sim \mathbf{p}_\theta(\widehat{p}_n\,|\widehat{p}_{n-1}, Y_n, Y_{n-1}\,) g_{\theta^*}(Y_n\,|X_n\,) f_{\theta^*}(X_n\,|X_{n-1}\,). \quad (18)$$

Using the deterministic mappings in (4) and (5) it can be easily shown that $\mathbf{p}_\theta(\widehat{p}_n\,|\widehat{p}_{n-1}, Y_n, Y_{n-1}\,)$ is equal to

$$\mathbf{p}_\theta(\widetilde{X}_n^{(1:N)}, X_{n-1}^{(1:N)}\,\Big|\,\widetilde{X}_{n-1}^{(1:N)}, X_{n-2}^{(1:N)}, Y_n, Y_{n-1}\,) = Q_\theta\left(\widetilde{X}_n^{(1:N)}\,\Big|\,X_{n-1}^{(1:N)}, Y_n\right) L\left(I_{n-1}^{(1:N)}|\widetilde{a}_{\theta,n-1}^{(1:N)}\right). \quad (19)$$

Sampling from (19) is equivalent to running the generic SMC algorithm from time $n-1$ to n. Equation (19) implies that the score estimate in (13) specifies to

$$\sum_{m=0}^{n}\left[\frac{\nabla_\theta Q_\theta\left(\widetilde{X}_m^{(1:N)}|X_{m-1}^{(1:N)}, Y_m\right)}{Q_\theta\left(\widetilde{X}_m^{(1:N)}|X_{m-1}^{(1:N)}, Y_m\right)} + \frac{\nabla_\theta L\left(I_{m-1}^{(1:N)}|\widetilde{a}_{\theta,m-1}^{(1:N)}\right)}{L\left(I_{m-1}^{(1:N)}|\widetilde{a}_{\theta,m-1}^{(1:N)}\right)}\right] \quad (20)$$

Note that because we only use a finite number N of particles, $\widehat{p}_n$ is only an approximation to the exact density $p_\theta\,(\,x_{n-1}|\,Y_{1:n-1})$ $\times q_\theta\,(\,x_n|\,Y_n, x_{n-1})$. Hence the maximizing value ϑ^* in (17) will not be exactly equal to the true parameter θ^*. In principle, as N increases, $J\,(\theta)$ will get closer to $l\,(\theta)$ and ϑ^* will converge to θ^*. Our simulation results indicate that ϑ^* provides a good approximation to θ^* for a moderate number of particles.

3.1. Variance reduction methods

Unlike the standard applications in the literature, we are dealing with a controlled Markov chain of a high dimension, because of the particle filter component of the state. This necessitates variance reductions methods for the gradient estimate. The score estimate $S_n\,(\theta)$ and consequently the gradient estimate $\widehat{\nabla J}_n(\theta)$ suffer from very large estimation variance that increases with time and the number of particles N. We therefore consider a number of variance reduction techniques:

1. Parallel Filters

The excessive variance of the gradient estimate is prohibitive for practical parameter estimation method within a reasonable number of iterations. We address this limitation by using P parallel particle filters. This effectively gives P independent draws from (19) that can be combined into an averaged unbiased estimate $\overline{\nabla J}_n(\theta_{n-1}) = \frac{1}{P}\sum_{j=1}^{P}\widehat{\nabla J}_{n,j}(\theta_{n-1})$. Note that in practice this also allows a lower number of particles to be used per particle filter. This approach can be viewed as having a single particle filter with $N \times P$ particles, where the particles are resampled within each of the P blocks and the blocks are not interacting.

2. Discount factor

The variance of the score function typically grows with time. We limit this variance increase at a cost of introducing some bias by using a discount factor $\varrho \in [0, 1)$ in the score iteration. For the values of ϱ used in our simulations, the bias was negligible.

3. Optimal Control Variate

Using the fact that $\mathbb{E}_\theta\,[S_n\,(\theta)] = 0$, we can achieve further variance reduction by considering a new estimate $\widetilde{\nabla J}_n(\theta) = \widehat{\nabla J}_n(\theta) - bS_n(\theta)$ for a suitable constant b. Subtracting $bS_n(\theta)$ does not add any bias but reduces the variance provided b is chosen properly. The optimal value for b, denoted b^*, is the solution of a quadratic cost function, namely the variance of $\widetilde{\nabla J}_n(\theta)$, and can itself be solved for using a second SA algorithm. This leads to the following two-time scale SA algorithm,

$$b_n = b_{n-1} - \zeta_n\left[b_{n-1}S_n^2\,(\theta_{n-1}) - \widehat{\nabla J}_n(\theta_{n-1})S_n(\theta_{n-1})\right],$$
$$\theta_n = \theta_{n-1} + \gamma_n\left[\widehat{\nabla J}_n(\theta_{n-1}) - b_n S_n(\theta_{n-1})\right]. \quad (21)$$

Here, the step-size $\{\zeta_n\}$ tends to zero more slowly than $\{\gamma_n\}$ does. For more details and exact conditions on the step sizes see [7]. In the implementation of the algorithm, we however use constant step-sizes. Note that since the parallel implementation provides us with averaged estimates, we use $\overline{S_n}(\theta_{n-1}) = \frac{1}{P}\sum_{j=1}^{P} S_{n,j}(\theta_{n-1})$ and $\overline{\nabla J}_n(\theta_{n-1})$ in (21) instead.

At time $n-1$, let $[\widehat{p}_{n-1,1}, \ldots, \widehat{p}_{n-1,P}]$ be P independent particle filters with corresponding score estimates $[S_{n-1,1}\,(\theta_{n-1}), \ldots,$ $S_{n-1,P}\,(\theta_{n-1})]$. The full algorithm for RML parameter estimation with variance reduction proceeds as follows:

Generic Recursive Maximum Likelihood algorithm

- (X_n, Y_n) is generated by the true system and only Y_n is observed
- For $j = 1, ..., P$
  - generate $\widehat{p}_{n,j} \sim p_{\theta_{n-1}}\,(\cdot\,|\widehat{p}_{n-1,j}, Y_{n:n-1}\,)$ using (19), i.e. run the SMC algorithm of section 1.2.
  - evaluate

$\widehat{\nabla J}_{n,j}(\theta_{n-1}) = \nabla\varphi_{\theta_{n-1}}(Y_n, \widehat{p}_{n,j}) + \varphi_{\theta_{n-1}}(Y_n, \widehat{p}_{n,j})S_{n,j}(\theta_{n-1})$
using $S_{n,j}(\theta) = \varrho S_{n-1,j}(\theta) + \frac{\nabla \mathbf{p}_\theta(\widehat{p}_{n,j}|Y_{n:n-1}, \widehat{p}_{n-1,j}\,)}{\mathbf{p}_\theta(\widehat{p}_{n,j}|Y_{n:n-1}, \widehat{p}_{n-1,j}\,)}$

- Calculate the averages $\overline{S_n}(\theta_{n-1}) = \frac{1}{P}\sum_{j=1}^{P} S_{n,j}(\theta_{n-1})$ and $\overline{\nabla J}_n(\theta_{n-1}) = \frac{1}{P}\sum_{j=1}^{P}\widehat{\nabla J}_{n,j}(\theta_{n-1})$
- Update control variate constant

$b_n = b_{n-1} - \zeta_n\left[b_{n-1}\overline{S_n}^2\,(\theta_{n-1}) - \overline{\nabla J}_n(\theta_{n-1})\overline{S_n}(\theta_{n-1})\right]$

- Calculate $\widetilde{\nabla J}_n(\theta_{n-1}) = \overline{\nabla J}_n(\theta_{n-1}) - b_n\overline{S_n}(\theta_{n-1})$
- Update parameter $\theta_n = \theta_{n-1} + \gamma_n\widetilde{\nabla J}_n(\theta_{n-1})$

Remark 1 *One draw of $Z_n = (X_n, Y_n, \widehat{p}_n)$ from (18) amounts to one-step evolution of the unknown true system in (6)-(7) and one iteration of the SMC algorithm in section 1.2. Here instead of a single SMC iteration, we use P parallel SMC iterations and obtain an average estimate to reduce the high variance.*

4. APPLICATIONS

4.1. Linear Gaussian state space model

Consider the following Linear Gaussian state space model

$$X_n = \phi X_{n-1} + \sigma_v V_n,\ X_0 \sim \mathcal{N}\left(0, \frac{\sigma_v^2}{1-\phi^2}\right)$$

$$Y_n = X_n + \sigma_w W_n$$

where $V_n \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0,1)$, $W_n \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0,1)$ and $\theta \triangleq (\sigma_v, \phi, \sigma_w)$ is unknown. We use the optimal importance density $q_\theta(X_n|Y_n, X_{n-1})$ and set the true parameter to $\theta^* \triangleq (0.2, 0.9, 0.15)$. The RML algorithm was implemented using $N = 500$ particles, $P = 50$ parallel filters, a discount factor of $\varrho = 0.5$ and the parameter values were randomly initialized. Figure 1 displays the MLE estimates and the associate control variate constants.

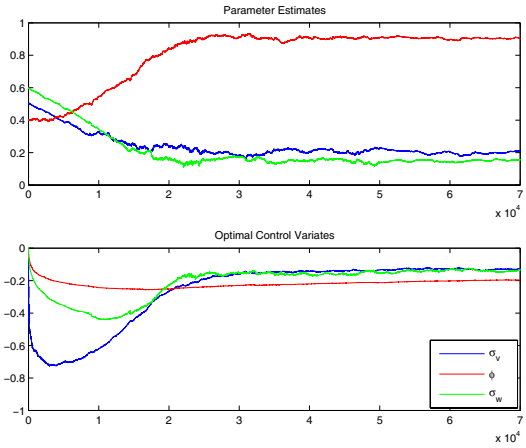

Fig. 1. On-line control variate constants and parameter estimates $\theta_n = [\sigma_{v,n}, \phi_n, \sigma_{w,n}]$ for the linear Gaussian model with true parameter $\theta^* = [0.2, 0.9, 0.15]$. Parameter estimates from top to bottom: ϕ_n, $\sigma_{v,n}$ and $\sigma_{w,n}$.

4.2. Stochastic Volatility model

The Stochastic Volatility model

$$X_n = \phi X_{n-1} + \sigma V_n,\ X_0 \sim \mathcal{N}\left(0, \frac{\sigma^2}{1-\phi^2}\right)$$

$$Y_n = \beta \exp\left(\frac{X_n}{2}\right) W_n$$

was considered, with $V_n \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0,1)$ and $W_n \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0,1)$. We are interested in estimating the parameter $\theta \triangleq (\phi, \sigma_V, \beta)$. The true parameter values were set to $\theta^* \triangleq (0.6, 0.9, 0.7)$. The RML algorithm was implemented using the prior as the importance density, i.e. $q_\theta(X_n|Y_n, X_{n-1}) = p_\theta(X_n|X_{n-1})$. An example of the estimates obtained using $N = 500$ particles, $P = 50$ parallel filters and discount factor of $\varrho = 0.2$ is shown in Figure 2.

5. DISCUSSION

In this paper we have proposed new particle methods to perform on-line maximum likelihood parameter estimation in general state space models. Our method is based on the key fact that the hidden state, its observation and the particle filter form an ergodic Markov chain that is controlled by the unknown static parameter of the hidden Markov model. Maximum likelihood estimates are obtained by optimizing a suitable average reward function using a score function approach and SA techniques. A batch implementation of the proposed algorithm can also be obtained.

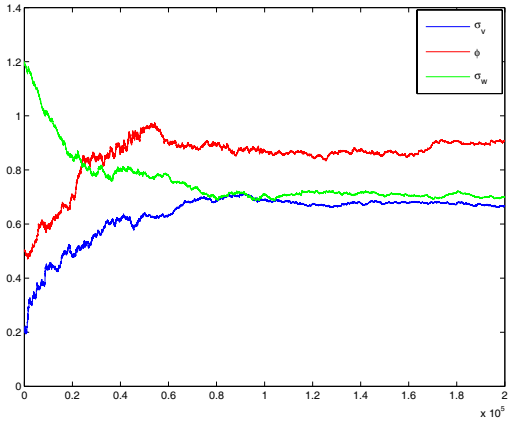

Fig. 2. On-line parameter estimates $\theta_n = [\sigma_n, \phi_n, \beta_n]$ for the Stochastic Volatility model with true parameter $\theta^* = [0.6, 0.9, 0.7]$. From top to bottom: ϕ_n, β_n and σ_n.

6. REFERENCES

[1] Andrieu C., Doucet A. and Tadic V.B. (2005) Online parameter estimation in general state space models, *Proc. IEEE CDC/ECC*

[2] Doucet A. and Tadić V. B. (2002) On-line Optimization of Sequential Monte Carlo Methods using Stochastic Approximation, *Proc. American Control Conference*.

[3] Doucet A. and Tadić V.B. (2003) Parameter Estimation in General State-Space Models using Particle Methods, *Ann. Inst. Stat. Math.*, vol. 55, no. 2, pp. 409-422.

[4] Doucet A., Godsill S.J. & Andrieu C. (2000) On sequential Monte Carlo sampling methods for Bayesian filtering. *Statist. Comput.*, **10**, 197-208.

[5] Fearnhead P. (2002) MCMC, sufficient statistics and particle filter, *J. Comp. Graph. Stat.*, vol. 11, pp. 848-862.

[6] Glynn, P. W. (1990), Likelihood Ratio Gradient Estimation for Stochastic Systems, *Communications of the ACM* 33, 75-84.

[7] Konda V. R. and Tsitsiklis J. N. (2003). Linear Stochastic Approximation Driven by Slowly Varying Markov Chains. *Systems and Control Letters*, Vol. 50, No. 2, pp. 95-102.

[8] LeGland F. and Mevel, L. (1997) Recursive estimation of hidden Markov models, *Proc. 36th IEEE Conf. Decision Control*.

[9] Le Gland F. and Mevel L. (1997). Asymptotic Properties of the MLE in HMM's. Proceedings of the 4th ECC.

[10] Liu J. and West M. (2001) Combined parameter and state estimation in simulation-based filtering, In *Sequential Monte Carlo Methods in Practice* (eds Doucet A., de Freitas J.F.G. and Gordon) N.J. New York: Springer-Verlag.

[11] Pflug G.C. (1996) *Optimization of Stochastic Models*. Kluwer.

[12] Poyiadjis G., Doucet A. and Singh S.S. (2005) Particle methods for optimal filter derivative: Application to parameter estimation, *Proceedings IEEE ICASSP*.

[13] Storvik G. (2002) Particle filters in state space models with the presence of unknown static parameters. *IEEE Trans. Signal Processing*, 50, 281-289.

[14] Tadić V.B. and Doucet A. (2002) Exponential forgetting and geometric ergodicity in state-space models. *Proceedings 41th IEEE Conf. Decision and Control*.