Proceeding of the 2004 American Control Conference
Boston, Massachusetts June 30 - July 2, 2004

FrP07.5

A Simulation Based Algorithm for Optimal Quantization in Non-Linear/Non-Gaussian State-Space Models

Vladislav B. Tadić and Arnaud Doucet

***Abstract*— The problem of the approximation of the optimal filter for non-linear/non-Gaussian state-space models is considered. This problem is studied for models with a multi-dimensional (continuous) state space and one-dimensional (continuous) observation space. An approximation of the optimal filter based on quantization is proposed. We quantize both the state and observation processes to obtain a hidden Markov model with discrete state and observation spaces for which the optimal filter can be computed exactly. The problem of the optimal selection of the parameters of this approximating model (quantization thresholds, states, transition probabilities, likelihood probabilities) is considered. An algorithm based on Monte Carlo gradient estimation and stochastic approximation is proposed. The asymptotic properties of the proposed algorithm are analyzed and sufficient conditions for its convergence are obtained.**

***Index Terms*— Quantization, hidden Markov models, optimal filtering, Monte Carlo gradient estimation, stochastic approximation.**

I. Introduction

Non-linear/non-Gaussian state-space models (also known as hidden Markov models with continuous state and observation spaces) are a broad class of stochastic processes capable of modeling very complex real-world dynamic systems. They have found a wide range of application in the areas such as automatic control, signal processing, machine learning, bioinformatics (just to name a few). Related to non-linear/non-Gaussian state-space models, one of the most important problems is the optimal filtering, i.e., the optimal estimation of their states in a situation when only their imperfect and noisy observations are available. Unfortunately, except in a few special cases, including linear Gaussian state-space models and hidden finite state space Markov chains, the optimal filters do not admit a finitely dimensional representation. However, nowadays the complexity of real-world systems coupled to the requirement of high performance precludes these simplifying assumptions. Moreover, in applications where bandwidth is scarce, it is necessary to minimize the number of bits used in representing the observation process and the optimal filter.

In this paper, the problem of the approximation of the optimal filter for non-linear/non-Gaussian state-space models is considered. This problem is studied for the models with a multi-dimensional (continuous) state space and one-dimensional (continuous) observation space, and an approximation of the optimal filter based on optimal quantization is proposed. We approximate the non-linear non-Gaussian state-space model by a hidden Markov model with discrete state and observation spaces. This approximating model allows to compute the optimal filter exactly while allowing to control the number of bits used in representing the observation process and the optimal filter. The problem of the optimal selection of the parameters of this approximating model (quantization thresholds, states, transition probabilities, likelihood probabilities) is considered. The optimality criterion is the mean square error between the true state process and the conditional expectation computed through the optimal filter associated to the approximating model. Minimizing this cost function is however complex as it involves the invariant distributions of both the original state-space model and of the filter of the approximating model. An algorithm based on Monte Carlo gradient estimation and stochastic approximation is proposed. The asymptotic properties of the proposed algorithm are analyzed and sufficient conditions for its convergence are obtained. (Simulation results will be presented in the last version of the paper.)

To the best of our knowledge, the quantization of hidden Markov models and the quantization based approximation of their optimal filters have only been studied in [6], [7], [9]. The approximation of the optimal filter presented in [6] is only based on the quantization (i.e., discretization) of the state space and is computationally more complex than the one proposed in this paper. Moreover, this quantization is not performed so as to optimize the performance of the filter associated to the quantized model. In [7], [9], only hidden Markov models with a discrete state-space have been considered.

The paper is organized as follows. In Section II, the signal model and the approximation of its optimal filter are defined. A criterion function for the optimal selection of the approximation parameters is also defined in Section II. In Section III, Monte Carlo gradient estimates of the criterion function are derived, while the algorithm for the criterion function minimization is proposed in Section IV. Results on the asymptotic behavior of the proposed algorithm are presented in Section V.

II. Signal Model and Approximation of the Optimal Filter

In this paper, the problems of the quantization of hidden Markov models with continuous state and observation spaces (i.e, of non-linear/non-Gaussian state-space models) and the approximation of their optimal filters is studied.

Vladislav B. Tadić is with Department of Automatic Control and Systems Engineering, University of Sheffield, Sheffield S1 3JD, United Kingdom, e-mail: v.tadic@sheffield.ac.uk

Arnaud Doucet is with Department of Engineering, Cambridge University, Cambridge CB2 1PZ, United Kingdom, e-mail: ad2@eng.cam.uk.ac

Hidden Markov models considered in this paper can be defined as a two component Markov chain $\{X_n, Y_n\}_{n\geq 0}$, where $\{X_n\}_{n\geq 0}$ and $\{Y_n\}_{n\geq 0}$ are R^{d_x}-valued and R-valued (respectively) stochastic processes which are defined on a probability space $(\Omega, \mathcal{F}, \mathcal{P})$ and satisfy the following relations:

$$\begin{aligned}&\mathcal{P}(X_{n+1}\in B|X^n, Y^n)\\ &= P(X_n, B) \ \ w.p.1, \quad \forall B\in\mathcal{B}^{d_x}, \ n\geq 0,\end{aligned}$$

$$\begin{aligned}&\mathcal{P}(Y_{n+1}\in B|X^{n+1}, Y^n)\\ &= \int_B q(X_{n+1}, y)dy \ \ w.p.1, \quad \forall B\in\mathcal{B}, \ n\geq 0.\end{aligned}$$

$X^n = [X_0^T \cdots X_n^T]^T$ and $Y^n = [Y_0 \cdots Y_n]^T$, $n \geq 0$, while $\mathcal{B}^{d_x}$, $\mathcal{B}$ are the families of Borel sets from R^{d_x}, R (respectively). $P(x, \cdot)$, $x\in R^{d_x}$, is a transition probability kernel (i.e., $P(x,\cdot)$ is a probability measure for all $x\in R^{d_x}$, and $P(\cdot, B)$ is Borel-measurable for all $B\in\mathcal{B}^{d_x}$), while $q: R^{d_x}\times R\to[0,\infty)$ is a Borel-measurable function with the following property:

$$\int q(x,y)dy = 1, \quad \forall x\in R^{d_x}.$$

The process $\{X_n\}_{n\geq 0}$ is observable only through the process $\{Y_n\}_{n\geq 0}$, i.e., the only information on $\{X_n\}_{n\geq 0}$ available to us is contained in $\{Y_n\}_{n\geq 0}$.

The problem we consider in this paper is the quantization of the process $\{Y_n\}_{n\geq 0}$ and the complexity reduction of the optimal filter $E(X_n|Y^n)$, $n\geq 0$. We want to represent (i.e., to quantize) $\{Y_n\}_{n\geq 0}$ by a signal taking N_y values ($N_y\geq 2$), and then, using this representation (i.e., quantized signal), to estimate $\{X_n\}_{n\geq 0}$ by the optimal filter for a hidden Markov model with N_x R^{d_x}-valued states ($N_x\geq 2$) and N_y observations (i.e., to approximate $E(X_n|Y^n)$, $n\geq 0$, with the optimal filter for the previously described hidden Markov model). Moreover, given N_x, N_y, we want to find the parameters of the quantizer (quantization thresholds) and the approximating hidden Markov model (states, transition probabilities and likelihood probabilities) which are optimal in the mean-square sense, i.e., which minimize the mean-square error of the estimates of $\{X_n\}_{n\geq 0}$ (obtained in the way described above). In order to achieve these objectives, the quantization thresholds and the states of the approximating hidden Markov model are considered as unknown parameters, while an additional parameterization is introduced to represent the transition and likelihood probabilities of the approximating hidden Markov model.

Let $\Theta = R^{N_x^2+N_xN_y+N_xd_x+1}\times[0,\infty)^{N_y-2}$. For $u_i = [u_{i,1}\cdots u_{i,N_x}]^T \in R^{N_x}$, $v_i = [v_{i,1}\cdots v_{i,N_y}]^T\in R^{N_y}$, $a_i = [a_{i,1}\cdots a_{i,d_x}]^T \in R^{d_x}$, $1\leq i\leq N_x$, $b = [b_1\cdots b_{N_y-1}]^T \in R\times[0,\infty)^{N_y-2}$, and $\theta = [u_1^T\cdots u_{N_x}^T \ v_1^T\cdots v_{N_x}^T \ a_1^T\cdots a_{N_x}^T \ b^T]^T$, let $u_i(\theta) = u_i$, $u_{i,j}(\theta) = u_{i,j}$, $v_i(\theta) = v_i$, $v_{i,k}(\theta) = v_{i,k}$, $a_i(\theta) = a_i$, $a_{i,l}(\theta) = a_{i,l}$, $b_m(\theta) = b_m$, $1\leq i,j\leq N_x$, $1\leq k\leq N_y$, $1\leq l\leq d_x$, $1\leq m< N_y$, $b(\theta) = b$, $c_0(\theta) = -\infty$, $c_{N_y}(\theta) = \infty$ and

$$\hat{p}_{i,j}(\theta) = \frac{\exp(u_{i,j})}{\sum_{k=1}^{N_x}\exp(u_{i,k})}, \quad 1\leq i,j\leq N_x,$$

$$\hat{q}_{i,j}(\theta) = \frac{\exp(v_{i,j})}{\sum_{k=1}^{N_y}\exp(v_{i,k})}, \quad 1\leq i\leq N_x, \ 1\leq j\leq N_y,$$

$$A(\theta) = [a_1\cdots a_{N_x}],$$

$$c_j(\theta) = \sum_{i=1}^{j} b_i, \quad 1\leq j< N_y,$$

while

$$\hat{P}(\theta) = [\hat{p}_{i,j}(\theta)]_{1\leq i,j\leq N_x},$$

$$\hat{Q}(\theta, i) = \text{diag}\{\hat{q}_{1,i}(\theta),\ldots,\hat{q}_{N_x,i}(\theta)\}, \quad 1\leq i\leq N_y.$$

$$\mathcal{Q}(\theta, y) = \sum_{i=1}^{N_y} iI_{(c_{i-1}(\theta), c_i(\theta)]}(y), \quad y\in R.$$

For any real-valued sequence $\{t_n\}_{n\geq 0}$, let $t^n = [t_0\cdots t_n]^T$. For $\theta\in\Theta$, $\nu\in[0,\infty)^{N_x}$ and any sequence $\{k_n\}_{n\geq 0}$ from $\{1,\ldots,N_y\}$, let $\nu_0(\theta, k^0, \nu) = \nu$ and

$$\mu_{n+1}(\theta, k^n, \nu) = \hat{P}^T(\theta)\nu_n(\theta, k^n, \nu), \quad n\geq 0,$$

$$\nu_{n+1}(\theta, k^{n+1}, \nu) = \frac{\hat{Q}(\theta, k_{n+1})\mu_{n+1}(\theta, k^n, \nu)}{e^T\hat{Q}(\theta, k_{n+1})\mu_{n+1}(\theta, k^n, \nu)}, \quad n\geq 0,$$

$$\phi_n(\theta, k^n, \nu) = A(\theta)\nu_n(\theta, k^n, \nu), \quad n\geq 0,$$

where $e = [1\cdots 1]^T\in R^{N_x}$. For $\theta\in\Theta$, $\nu\in[0,\infty)^{N_x}$, let

$$\hat{Y}_n(\theta) = \mathcal{Q}(\theta, Y_n), \quad n\geq 0,$$

$$\hat{X}_n(\theta, \nu) = \phi_n(\theta, \hat{Y}^n(\theta), \nu), \quad n\geq 0,$$

while $\{\tilde{X}_n(\theta)\}_{n\geq 0}$ and $\{\tilde{Y}_n(\theta)\}_{n\geq 0}$ are $\{a_1(\theta),\ldots,a_{N_x}(\theta)\}$-valued and $\{1,\ldots,N_y\}$-valued stochastic processes satisfying

$$\begin{aligned}&\mathcal{P}(\tilde{X}_{n+1}(\theta) = a_j|\tilde{X}^n(\theta), \tilde{Y}^n(\theta))\\ &= \sum_{i=1}^{N_x}\hat{p}_{i,j}(\theta)I_{\{a_i(\theta)\}}(\tilde{X}_n(\theta)) \ \ w.p.1, \ 1\leq j\leq N_x,\end{aligned}$$

$$\begin{aligned}&\mathcal{P}(\tilde{Y}_{n+1}(\theta) = j|\tilde{X}^{n+1}(\theta), \tilde{Y}^n(\theta))\\ &= \sum_{i=1}^{N_x}\hat{q}_{i,j}(\theta)I_{\{a_i(\theta)\}}(\tilde{X}_{n+1}(\theta)) \ \ w.p.1, \ 1\leq j\leq N_y.\end{aligned}$$

Then, it is straightforward to verify that $\phi_n(\theta, \tilde{Y}^n(\theta), \nu)$, $n\geq 0$, is the optimal filter for the hidden Markov model $\{\tilde{X}_n(\theta), \tilde{Y}_n(\theta)\}_{n\geq 0}$, i.e,

$$E(\tilde{X}_n(\theta)|\tilde{Y}^n(\theta)) = \phi_n(\theta, \tilde{Y}^n(\theta), \nu), \quad n\geq 0,$$

if $\mathcal{P}(\tilde{X}_0(\theta) = i) = \nu_i$, $1\leq i\leq N_x$, where ν_i is the i-th component of ν. On the other hand, $\{\hat{Y}_n(\theta)\}_{n\geq 0}$ can be considered as the representation of $\{Y_n\}_{n\geq 0}$ which takes values $\{1,\ldots,N_y\}$, while $\mathcal{Q}(\theta,\cdot)$ is the quantizer for $\{Y_n\}_{n\geq 0}$. Moreover, $\phi_n(\theta, \hat{Y}^n(\theta), \nu)$, $n\geq 0$, can be thought of as an approximation of the optimal filter

$E(X_n|Y^n)$, $n \geq 0$, while $\{\hat{X}_n(\theta,\nu)\}_{n\geq 0}$ are the estimates of $\{X_n\}_{n\geq 0}$ produced by the approximating filter $\phi_n(\theta,\cdot,\nu)$, $n \geq 0$.

For $\theta \in \Theta$, $\nu \in [0,\infty)^{N_x}$, let

$$J_n(\theta,\nu) = 2^{-1}E\|X_n - \hat{X}_n(\theta,\nu)\|^2, \quad n \geq 0,$$

while $J(\theta) = \lim_{n\to\infty} J_n(\theta,\nu)$ (notice that $\lim_{n\to\infty} J_n(\theta,\nu)$ is well-defined, finite and does not depend on ν if $\{X_n\}_{n\geq 0}$ is geometrically ergodic and the transition probability matrix of $\{\tilde{X}_n(\theta)\}_{n\geq 0}$ is strictly positive; for details see [5], [8]). Our objective is to determine $\theta \in \Theta$ which provides the mean-square optimal estimation of $\{X_n\}_{n\geq 0}$ by $\{\hat{X}_n(\theta,\nu)\}_{n\geq 0}$, i.e., which minimizes the asymptotic mean-square error $J(\cdot)$. Once the optimal $\theta \in \Theta$ is determined, the corresponding parameters of the quantizer $\mathcal{Q}(\theta,\cdot)$ (the threshold levels $c_k(\theta)$, $1 \leq k < N_y$) and the approximating hidden Markov chain $\{\tilde{X}_n(\theta), \tilde{Y}_n(\theta)\}_{n\geq 0}$ (the transition probability matrix $\hat{P}(\theta)$, the likelihood probabilities $\hat{Q}(\theta,k)$, $1 \leq k \leq N_y$, the states $a_1(\theta),\ldots,a_{N_x}(\theta)$) can easily be computed. We study the problem of the minimization of $J(\cdot)$ for the case where $P(\cdot,\cdot)$, $q(\cdot,\cdot)$ are known and easy to be sampled from, and where $\int_{-\infty}^{y} q(x,y')dy'$ can analytically be computed for all $x \in R^{d_x}$, $y \in R$.

The rationale for quantizing the observations $\{Y_n\}_{n\geq 0}$ and approximating the optimal filter $E(X_n|Y^n)$, $n \geq 0$, with $\phi_n(\theta,\hat{Y}^n(\theta),\nu)$, $n \geq 0$, comes out from the following fact: For optimal $\theta \in \Theta$, the mean-square error of the approximate estimates $\hat{X}_n(\theta,\nu)$, $n \geq 0$, is close to the mean-square error of the optimal estimates $E(X_n|Y^n)$, $n \geq 0$, if N_x, N_y are sufficiently large, i.e., $\min_{\theta\in\Theta} J(\theta) \approx E\|X_n - E(X_n|Y^n)\|^2$ for $N_x, N_y \gg 1$ (for details see [10]).

III. Gradient Estimates of the Mean-Square Error

In this section, we present unbiased (Monte Carlo) estimators for the gradients of $J(\cdot)$ and $J_n(\cdot,\nu)$, $\nu \in [0,\infty)$, $n \geq 0$ (for the derivation see [10]). The reason for using Monte Carlo simulations for estimating the gradients of $J(\cdot)$ comes out from the following fact: $J(\cdot)$ does not admit a closed form expression (as it depends on invariant distribution of the augmented Markov chain $\{X_n, \hat{Y}_n(\theta), \nu_n(\theta,\hat{Y}^n(\theta),\nu)\}_{n\geq 0}$, which itself does not have a closed form solution), while Monte Carlo methods provide a computationally inexpensive and efficient way to estimate those derivatives.

For $\theta \in \Theta$, let

$$r(\theta,x,i) = \int_{c_{i-1}(\theta)}^{c_i(\theta)} q(x,y)dy, \quad x \in R^{d_x},\ 1 \leq i \leq N_y,$$

while $\hat{r}(\theta,x,i) \in R^{N_y-1}$, $1 \leq i \leq N_y$, are the vectors defined as

$$[\hat{r}(\theta,x,i)]_j = \begin{cases} q(x,c_i(\theta)) - q(x,c_{i-1}(\theta)), & \text{if } i > j \\ q(x,c_i(\theta)), & \text{if } i = j \\ 0, & \text{otherwise} \end{cases},$$
$$x \in R^{d_x},\ 1 \leq i \leq N_y,\ 1 \leq j < N_y.$$

For $\theta \in \Theta$, let $\tilde{p}_{i,j}(\theta), \tilde{q}_{i,k}(\theta,l) \in R^{N_x}$, $1 \leq i,j \leq N_x$, $1 \leq k,l \leq N_y$, be the vectors defined as

$$[\tilde{p}_{i,j}(\theta)]_k = \begin{cases} \hat{p}_{i,j}(\theta)(1-\hat{p}_{i,j}(\theta)), & \text{if } j = k \\ -\hat{p}_{i,j}(\theta)\hat{p}_{i,k}(\theta), & \text{otherwise} \end{cases},$$
$$1 \leq i,j \leq N_x,$$

$$[\tilde{q}_{i,j}(\theta,k)]_l = \begin{cases} \hat{q}_{i,j}(\theta)(1-\hat{q}_{i,j}(\theta)), & \text{if } i = l, j = k \\ -\hat{q}_{i,k}(\theta)\hat{q}_{i,l}(\theta), & \text{if } i = l, j \neq k \\ 0, & \text{otherwise} \end{cases},$$
$$1 \leq i,l \leq N_x,\ 1 \leq j,k \leq N_y,$$

while

$$\tilde{P}_i(\theta) = [\tilde{p}_{i,1}(\theta) \cdots \tilde{p}_{i,N_x}(\theta)], \quad 1 \leq i \leq N_x,$$

$$\tilde{Q}_i(\theta,j) = [\tilde{q}_{i,1}(\theta,j) \cdots \tilde{q}_{i,N_y}(\theta,j)],$$
$$1 \leq i \leq N_x,\ \ 1 \leq j \leq N_y.$$

For $\theta \in \Theta$, let $\partial/\partial u_{i,j}(\theta)$, $\partial/\partial v_{i,k}(\theta)$, $\partial/\partial a_{i,l}(\theta)$, $\partial/\partial b_m(\theta)$, $1 \leq i,j \leq N_x$, $1 \leq k \leq N_y$, $1 \leq l \leq d_x$, $1 \leq m < N_y$, be the derivatives with respect to the corresponding components of θ (i.e., with respect to the $(N_x(j-1)+k)$-th, $(N_x^2+N_y(j-1)+l)$-th, $(N_x^2+N_xN_y+d_x(j-1)+m)$-th, $(N_x^2+N_xN_y+N_xd_x+j)$-th components of θ, respectively), while $\nabla_{u_i(\theta)} = [\partial/\partial u_{i,1}(\theta) \cdots \partial/\partial u_{i,N_x}(\theta)]^T$, $\nabla_{v_i(\theta)} = [\partial/\partial v_{i,1}(\theta) \cdots \partial/\partial v_{i,N_y}(\theta)]^T$, $\nabla_{a_i(\theta)} = [\partial/\partial a_{i,1}(\theta) \cdots \partial/\partial a_{i,d_x}(\theta)]^T$, $1 \leq i \leq N_x$, $\nabla_{b(\theta)} = [\partial/\partial b_1(\theta) \cdots \partial/\partial b_{N_y-1}(\theta)]^T$.

For $\theta \in \Theta$, $\nu \in [0,\infty)^{N_x}$ and any sequence $\{k_n\}_{n\geq 0}$ from $\{1,\ldots,N_y\}$, let

$$M'_{n+1,i}(\theta,k^n,\nu) = \left[\frac{\partial\mu_{n+1}(\theta,k^n,\nu)}{\partial u_{i,1}(\theta)} \cdots \frac{\partial\mu_{n+1}(\theta,k^n,\nu)}{\partial u_{i,N_x}(\theta)}\right],$$
$$1 \leq i \leq N_x,\ n \geq 0,$$

$$N'_{n,i}(\theta,k^n,\nu) = \left[\frac{\partial\nu_n(\theta,k^n,\nu)}{\partial u_{i,1}(\theta)} \cdots \frac{\partial\nu_n(\theta,k^n,\nu)}{\partial u_{i,N_x}(\theta)}\right],$$
$$1 \leq i \leq N_x,\ n \geq 0,$$

$$M''_{n+1,i}(\theta,k^n,\nu) = \left[\frac{\partial\mu_{n+1}(\theta,k^n,\nu)}{\partial v_{i,1}(\theta)} \cdots \frac{\partial\mu_{n+1}(\theta,k^n,\nu)}{\partial v_{i,N_y}(\theta)}\right],$$
$$1 \leq i \leq N_x,\ n \geq 0,$$

$$N''_{n,i}(\theta,k^n,\nu) = \left[\frac{\partial\nu_n(\theta,k^n,\nu)}{\partial v_{i,1}(\theta)} \cdots \frac{\partial\nu_n(\theta,k^n,\nu)}{\partial v_{i,N_y}(\theta)}\right],$$
$$1 \leq i \leq N_x,\ n \geq 0.$$

For $\theta \in \Theta$, $\nu \in [0,\infty)^{N_x}$, $\beta \in [0,1]$ and any sequences $\{x_n\}_{n\geq 0}$, $\{k_n\}_{n\geq 0}$ from R^{d_x}, $\{1,\dots,N_y\}$ (respectively), let

$$s_n^\beta(\theta, x^n, k^n) = \sum_{i=0}^{n} (1-\beta)^{n-i} \frac{\hat{r}(\theta, x_i, k_i)}{r(\theta, x_i, k_i)}, \quad n \geq 0, \quad (1)$$

while

$$\begin{aligned}\widehat{\nabla}_{u_i(\theta)} J_n(\theta,\nu) = &-(N'_{n,i}(\theta, \hat{Y}^n(\theta),\nu))^T A^T(\theta) \\ &\cdot (X_n - A(\theta)\nu_n(\theta, \hat{Y}^n(\theta),\nu)), \\ &1 \leq i \leq N_x, \ n \geq 0, \end{aligned} \quad (2)$$

$$\begin{aligned}\widehat{\nabla}_{v_i(\theta)} J_n(\theta,\nu) = &-(N''_{n,i}(\theta, \hat{Y}^n(\theta),\nu))^T A^T(\theta) \\ &\cdot (X_n - A(\theta)\nu_n(\theta, \hat{Y}^n(\theta),\nu)), \\ &1 \leq i \leq N_x, \ n \geq 0, \end{aligned} \quad (3)$$

$$\begin{aligned}\widehat{\nabla}_{a_i(\theta)} J_n(\theta,\nu) = &-(X_n - A(\theta)\nu_n(\theta, \hat{Y}^n(\theta),\nu)) \\ &\cdot \nu_{n,i}(\theta, \hat{Y}^n(\theta),\nu), \\ &1 \leq i \leq N_x, \ n \geq 0, \end{aligned} \quad (4)$$

$$\begin{aligned}\widehat{\nabla}_{b(\theta)} J_n^\beta(\theta,\nu) = &2^{-1} \|X_n - A(\theta)\nu_n(\theta, \hat{Y}^n(\theta),\nu)\|^2 \\ &\cdot s_n^\beta(\theta, X^n, \hat{Y}^n(\theta)), \\ &1 \leq i \leq N_x, \ n \geq 0. \end{aligned} \quad (5)$$

It is straightforward to verify that

$$N'_{0,i}(\theta, k^0, \nu) = 0 \in R^{N_x \times N_x}, \quad 1 \leq i \leq N_x, \quad (6)$$

$$N''_{0,i}(\theta, k^0, \nu) = 0 \in R^{N_x \times N_y}, \quad 1 \leq i \leq N_x, \quad (7)$$

$$\begin{aligned}M'_{n+1,i}(\theta, k^n, \nu) = &\hat{P}^T(\theta) N'_{n,i}(\theta, k^n, \nu) \\ &+ \tilde{P}_i(\theta)\nu_{n,i}(\theta, k^n, \nu), \\ &1 \leq i \leq N_x, \ n \geq 0, \end{aligned} \quad (8)$$

$$\begin{aligned}M''_{n+1,i}(\theta, k^n, \nu) = &\hat{P}^T(\theta) N''_{n,i}(\theta, k^n, \nu), \\ &1 \leq i \leq N_x, \ n \geq 0, \end{aligned} \quad (9)$$

$$\begin{aligned}N'_{n+1,i}(\theta, k^{n+1}, \nu) = &\frac{I - \nu_{n+1}(\theta, k^{n+1}, \nu) e^T}{e^T \hat{Q}(\theta, k_{n+1}) \mu_{n+1}(\theta, k^n, \nu)} \\ &\cdot \hat{Q}(\theta, k_{n+1}) M'_{n+1,i}(\theta, k^n, \nu), \\ &1 \leq i \leq N_x, \ n \geq 0, \end{aligned} \quad (10)$$

$$\begin{aligned}N''_{n+1,i}(\theta, k^{n+1}, \nu) = &\frac{I - \nu_{n+1}(\theta, k^{n+1}, \nu) e^T}{e^T \hat{Q}(\theta, k_{n+1}) \mu_{n+1}(\theta, k^n, \nu)} \\ &\cdot (\hat{Q}(\theta, k_{n+1}) M''_{n+1}(\theta, k^n, \nu) \\ &+ \tilde{Q}_i(\theta, k_{n+1}) \mu_{n+1,i}(\theta, k^n, \nu)), \\ &1 \leq i \leq N_x, \ n \geq 0, \end{aligned} \quad (11)$$

for all $\theta \in \Theta$, $\nu \in [0,\infty)^{N_x}$ and any sequence $\{k_n\}_{n\geq 0}$ from $\{1,\dots,N_y\}$, where $\mu_{n+1,i}(\theta, k^n, \nu)$, $\nu_{n,i}(\theta, k^n, \nu)$ are the i-th components of $\mu_{n+1}(\theta, k^n, \nu)$, $\nu_n(\theta, k^n, \nu)$ (respectively).

It can be demonstrated that $\widehat{\nabla}_{u_i(\theta)} J_n(\theta,\nu)$, $\widehat{\nabla}_{v_i(\theta)} J_n(\theta,\nu)$, $\widehat{\nabla}_{a_i(\theta)} J_n(\theta,\nu)$, $\widehat{\nabla}_{b(\theta)} J_n^0(\theta,\nu)$ are the unbiased estimators of $\nabla_{u_i(\theta)} J_n(\theta,\nu)$, $\nabla_{v_i(\theta)} J_n(\theta,\nu)$, $\nabla_{a_i(\theta)} J_n(\theta,\nu)$, $\nabla_{b(\theta)} J_n(\theta,\nu)$ (respectively) for all $\theta \in \Theta$, $\nu \in [0,\infty)^{N_x}$, $1 \leq i \leq N_x$, $n \geq 0$, i.e.,

$$\nabla_{u_i(\theta)} J_n(\theta,\nu) = E(\widehat{\nabla}_{u_i(\theta)} J_n(\theta,\nu)), \ 1 \leq i \leq N_x, n \geq 0,$$

$$\nabla_{v_i(\theta)} J_n(\theta,\nu) = E(\widehat{\nabla}_{v_i(\theta)} J_n(\theta,\nu)), \ 1 \leq i \leq N_x, n \geq 0,$$

$$\nabla_{a_i(\theta)} J_n(\theta,\nu) = E(\widehat{\nabla}_{a_i(\theta)} J_n(\theta,\nu)), \ 1 \leq i \leq N_x, n \geq 0,$$

$$\nabla_{b(\theta)} J_n(\theta,\nu) = E(\widehat{\nabla}_{b(\theta)} J_n(\theta,\nu)), \ n \geq 0,$$

for all $\theta \in \Theta$, $\nu \in [0,\infty)^{N_x}$. Moreover, the results of [5], [8] suggest that if $\{X_n\}_{n\geq 0}$ is geometrically ergodic, $\widehat{\nabla}_{u_i(\theta)} J_n(\theta,\nu)$, $\widehat{\nabla}_{v_i(\theta)} J_n(\theta,\nu)$, $\widehat{\nabla}_{a_i(\theta)} J_n(\theta,\nu)$, $\widehat{\nabla}_{b(\theta)} J_n^\beta(\theta,\nu)$ are the asymptotically unbiased estimators of $\nabla_{u_i(\theta)} J(\theta)$, $\nabla_{v_i(\theta)} J(\theta)$, $\nabla_{a_i(\theta)} J(\theta)$, $\nabla_{b(\theta)} J(\theta)$ (respectively) for all $\theta \in \Theta$, $\nu \in [0,\infty)^{N_x}$, $1 \leq i \leq N_x$, as $n \to \infty$, $\beta \to 0$.

IV. Simulation Based Optimization of the Mean-Square Error

In this section, we present an algorithm for finding the minima of $J(\cdot)$. This algorithm is based on the Monte Carlo estimators of the gradients of $J(\cdot)$ (derived in the previous section) and stochastic approximation (which uses those gradient estimates to search for minima of $J(\cdot)$).

The gradient estimates (2) – (5) (together with the supporting equations (6) – (11), (1)) and stochastic approximation suggest the following algorithm for minimizing $J(\cdot)$:

$$X_{n+1} \sim P(X_n, \cdot), \quad n \geq 0, \quad (12)$$

$$Y_{n+1} \sim q(X_{n+1}, \cdot), \quad n \geq 0, \quad (13)$$

$$\begin{aligned}\theta_n = &[u_{n,1}^T \cdots u_{n,N_x}^T v_{n,1}^T \cdots v_{n,N_x}^T \\ &a_{n,1}^T \cdots a_{n,N_x}^T b_n^T]^T, \quad n \geq 0, \end{aligned} \quad (14)$$

$$Z_{n+1} = \mathcal{Q}(\theta_n, Y_{n+1}), \quad n \geq 0, \quad (15)$$

$$\mu_{n+1} = \hat{P}^T(\theta_n)\mu_n, \quad n \geq 0, \quad (16)$$

$$\nu_{n+1} = \frac{\hat{Q}(\theta_n, Z_{n+1})\mu_{n+1}}{e^T \hat{Q}(\theta_n, Z_{n+1})\mu_{n+1}}, \quad n \geq 0, \quad (17)$$

$$\begin{aligned}M'_{n+1,i} = &\hat{P}^T(\theta_n) N'_{n,i} + \tilde{P}_i(\theta_n)\nu_{n,i}, \\ &1 \leq i \leq N_x, \ n \geq 0, \end{aligned} \quad (18)$$

$$M''_{n+1,i} = \hat{P}^T(\theta_n) N''_{n,i}, \quad 1 \leq i \leq N_x, \ n \geq 0, \quad (19)$$

$$\begin{aligned}N'_{n+1,i} = &\frac{I - \nu_{n+1} e^T}{e^T \hat{Q}(\theta_n, Z_{n+1})\mu_{n+1}} \hat{Q}(\theta_n, Z_{n+1}) M'_{n+1,i}, \\ &1 \leq i \leq N_x, \ n \geq 0, \end{aligned} \quad (20)$$

$$\begin{aligned}N''_{n+1,i} = &\frac{I - \nu_{n+1} e^T}{e^T \hat{Q}(\theta_n, Z_{n+1})\mu_{n+1}} \\ &\cdot (\hat{Q}(\theta_n, Z_{n+1}) M''_{n+1,i} + \tilde{Q}_i(\theta_n, Z_{n+1})\mu_{n+1,i}), \\ &1 \leq i \leq N_x, \ n \geq 0, \end{aligned} \quad (21)$$

$$s_{n+1} = (1-\beta)s_n + \frac{\hat{r}(\theta_n, X_{n+1}, Z_{n+1})}{r(\theta_n, X_{n+1}, Z_{n+1})}, \quad n \geq 0, \quad (22)$$

$$\widehat{\nabla}_{u_i} J_{n+1} = -(N'_{n+1,i})^T A^T(\theta_n)(X_{n+1} - A(\theta_n)\nu_{n+1}), \quad 1 \le i \le N_x,\ n \ge 0, \tag{23}$$

$$\widehat{\nabla}_{v_i} J_{n+1} = -(N''_{n+1,i})^T A^T(\theta_n)(X_{n+1} - A(\theta_n)\nu_{n+1}), \quad 1 \le i \le N_x,\ n \ge 0, \tag{24}$$

$$\widehat{\nabla}_{a_i} J_{n+1} = -(X_{n+1} - A(\theta_n)\nu_{n+1})\nu_{n+1,i} \quad 1 \le i \le N_x,\ n \ge 0, \tag{25}$$

$$\widehat{\nabla}_b J_{n+1} = 2^{-1}\|X_{n+1} - A(\theta_n)\nu_{n+1}\|^2 s_{n+1}, \quad n \ge 0, \tag{26}$$

$$u_{n+1,i} = u_{n,i} - \alpha_{n+1}\widehat{\nabla}_{u_i} J_{n+1}, \quad 1 \le i \le N_x,\ n \ge 0, \tag{27}$$

$$v_{n+1,i} = v_{n,i} - \alpha_{n+1}\widehat{\nabla}_{v_i} J_{n+1}, \quad 1 \le i \le N_x,\ n \ge 0, \tag{28}$$

$$a_{n+1,i} = a_{n,i} - \alpha_{n+1}\widehat{\nabla}_{a_i} J_{n+1}, \quad 1 \le i \le N_x,\ n \ge 0, \tag{29}$$

$$b_{n+1} = P_{\Theta_b}(b_n - \alpha_{n+1}\widehat{\nabla}_b J_{n+1}), \quad n \ge 0. \tag{30}$$

$\{\alpha_n\}_{n\ge1}$ is a sequence of positive reals and represents the algorithm step-sizes. $\beta \in (0,1)$ is a constant and can be considered as an error-forgetting and stabilizing factor in the recursion (22). $P_{\Theta_b}(\cdot)$ is the projection on $\Theta_b = R \times [0,\infty)^{N_y-2}$, i.e.,

$$P_{\Theta_b}(b) = \arg\min_{b' \in \Theta_b} \|b - b'\|, \quad b \in R^{N_y-1}.$$

$N'_{0,i} = 0 \in R^{N_x \times N_x}$, $N''_{0,i} = 0 \in R^{N_x \times N_y}$, $1 \le i \le N_x$, $s_0 = 0 \in R^{N_y-1}$, while $u_{0,i} \in R^{N_x}$, $v_{0,i} \in R^{N_y}$, $a_{0,i} \in R^{d_x}$, $1 \le i \le N_x$, $b_0 \in B$, $\nu_0 \in [0,\infty)^{N_x}$ are deterministic variables and represent the algorithm initial values. $\mu_{n+1,i}$, $\nu_{n,i}$ are the i-th components of μ_{n+1}, ν_n (respectively), $1 \le i \le N_x$, $n \ge 0$, while $X_{n+1} \sim P(X_n,\cdot)$, $Y_{n+1} \sim q(X_{n+1},\cdot)$ denote sampling from $P(X_n,\cdot)$, $q(X_{n+1},\cdot)$ (respectively), $n \ge 0$.

Remark: It can easily be deduced that the algorithm (12) – (30) falls into the category of stochastic approximation algorithms. Moreover, the algorithm (12) – (30) is of the same form as the algorithms analyzed in [1].

V. Asymptotic Analysis

In this section, we present results on the asymptotic behavior of the algorithm (12) – (30). The obtained results confirm that the algorithm (12) – (30) minimizes $J(\cdot)$ over Θ.

The asymptotic analysis of the algorithm (12) – (30) is carried out for the case where the algorithm step-sizes are constant $\alpha_n = \alpha$, $n \ge 1$, and where the step-size α and the forgetting factor β tend to zero. In order to emphasize the fact that the asymptotic behavior of the iterates $\{\theta_n\}_{n\ge0}$ depends on α, β, we use the notation $\theta_n^{\alpha,\beta}$ in what follows.

For $\rho \in (0,\infty)$, let $B_\Theta^\rho = \{\theta \in \Theta : \|\theta\| \le \rho\}$. The algorithm (12) – (30) is analyzed under the following assumptions:

A1: For all $\rho \in [1,\infty)$, there exists a Borel-measurable function $\phi_\rho : R^{d_x} \to [1,\infty)$ such that

$$\phi_\rho(x) \ge \|x\|^8,$$

$$\|\hat{r}(\theta,x,i)\| \le \phi_\rho^{1/4}(x) r(\theta,x,i), \quad 1 \le i \le N_y,$$

$$\left\| \frac{\hat{r}(\theta',x,i)}{r(\theta',x,i)} - \frac{\hat{r}(\theta'',x,i)}{r(\theta'',x,i)} \right\| \le \phi_\rho^{1/4}(x)\|\theta' - \theta''\|, \quad 1 \le i \le N_y,$$

for all $\theta, \theta', \theta'' \in B_\Theta^\rho$, $x \in R^{d_x}$.

A2: $\{X_n\}_{n\ge0}$ has a unique invariant probability measure $\pi(\cdot)$. For all $\rho \in [1,\infty)$, there exist constants $r_\rho \in (0,1)$, $K_\rho \in [1,\infty)$ such that

$$\left| \int \psi(x')P^n(x,dx') - \int \psi(x')\pi(dx') \right| \le K_\rho r_\rho^n \phi_\rho(x), \quad n \ge 1,$$

for all $x \in R^{d_x}$ and any Borel-measurable function $\psi : R^{d_x} \to R$ satisfying $0 \le \psi(x) \le \phi_\rho(x)$ for all $x \in R^{d_x}$.

A1 corresponds to the growth rate and smoothness of the functions $r(\cdot,\cdot,i)$, $\hat{r}(\cdot,\cdot,i)$, $1 \le i \le N_y$. It holds if $q(\cdot,\cdot)$ is locally Lipschitz continuous. A2 is related to the stability properties of $\{X_n\}_{n\ge0}$. It requires $\{X_n\}_{n\ge0}$ to be ϕ_ρ-geometrically ergodic for all $\rho \in [1,\infty)$.

Let $P_\Theta(\cdot)$ be the projection on Θ, i.e.,

$$P_\Theta(\theta) = \arg\min_{\theta' \in \Theta} \|\theta - \theta'\|, \quad \theta \in R^{N_x^2 + N_xN_y + N_xd_x + N_y - 1}.$$

The main result on the asymptotic behavior of the algorithm (12) – (30) is contained in the next theorem:

Theorem 1: Let $\theta_0^{\alpha,\beta} = \theta_0$ for all $\alpha, \beta \in (0,\infty)$, where $\theta_0 \in \Theta$ is a deterministic variable not depending on α, β. Moreover, let $\bar{\theta}_0^\alpha = \theta_0$ and

$$\bar{\theta}_{n+1}^\alpha = P_\Theta(\bar{\theta}_n^\alpha - \alpha \nabla J(\bar{\theta}_n^\alpha)), \quad n \ge 0. \tag{31}$$

Then,

$$\lim_{\alpha,\beta \to 0} \mathcal{P}\left(\sup_{0 \le n \le t/\alpha} \|\theta_n^{\alpha,\beta} - \bar{\theta}_n^\alpha\| \ge \delta \right) = 0$$

for all $\delta, t \in (0,\infty)$.

The proof is essentially based on the results of [1, Part II], [5], [8]. For details see [10].

Remark: Theorem 1 basically claims that the iterates $\{\theta_n^{\alpha,\beta}\}_{n\ge0}$ of the algorithm (12) – (30) asymptotically behave for $\alpha, \beta \to 0$ as the gradient search (31). Then, the further asymptotic properties of $\{\theta_n^{\alpha,\beta}\}_{n\ge0}$ can be inferred from the existing results on the (deterministic) projected gradient algorithms (see e.g., [2, Chapter 2] and references cited therein). The results of this type are typical for the asymptotic analysis of stochastic approximation algorithms (for details see [1], [4] and references cited therein).

VI. Application

We demonstrate our methodology on a non-linear state-space model. Let us consider the following stochastic volatility model arising in econometrics

$$X_{n+1} = \phi X_n + \sigma V_{n+1},$$
$$Y_n = \beta \exp(X_n/2) W_n,$$

where $V_n \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0,1)$ and $W_n \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0,1)$ are two mutually independent sequences of independent identically distributed (i.i.d.) Gaussian random variables, independent of the initial state X_0. The parameters are given by $(\beta, \phi, \sigma) = (1, 0.8, 1)$.

For this non-linear state space model, it is impossible to compute the filter in closed form and our quantization algorithm was compared to a particle approximation of the optimal filter (using $N = 10000$ particles).

We limit ourselves here to the quantization of the hidden state space. We use respectively $M = 10$, 50 and 100 quantization levels. We simulate $P = 10$ realizations of $T = 200$ observations. Our performance measure is the average mean square error (AMSE) between the true state and the approximation of its conditional expectation.

The results are given in the table below.

	$M = 10$	$M = 50$	$M = 100$	Particle filter
AMSE	0.321	0.218	0.183	0.142

Remark: Notice that the computational complexity of the used particle filter is equivant to computational complexity corresponding to 10000 quantization levels.

REFERENCES

[1] A. Benveniste, M. Metivier, and P. Priouret, *Adaptive Algorithms and Stochastic Approximation*, Springer Verlag, 1990.

[2] D. P. Bertsekas, *Nonlinear Programming*, Athena Scientific, 1999, 2nd edition.

[3] A. Gersho and R. M. Gray, *Vector Quantization and Signal Compression*, Kluwer Academic Publishers, 1992.

[4] H. J. Kushner and G. G. Yin, *Stochastic Approximation Algorithms and Applications*, Springer Verlag, 1997.

[5] F. LeGland and L. Mevel, “Exponential Forgetting and Geometric Ergodicity in Hidden Markov Models,” *Mathematics of Control, Signals and Systems*, vol. 13, 63-93, 2000.

[6] G. Pages and H. Pham, *Optimal Quantization Methods for Nonlinear Filtering with Discrete-Time Observations*, Prepublication PMA-778, Laboratoire de Probabilites et Modeles aleatoires, Universit Paris 6.

[7] L. Shue, S. Dey, B.D.O. Anderson, and F.D. Bruyne, “On State-Estimation of a Two-State Hidden Markov Model with Quantization,” *IEEE Transactions on Signal Processing*, vol.49, pp. 202-208, 2001.

[8] V. B. Tadić and A. Doucet, “Exponential Forgetting and Geometric Ergodicity for Optimal Filtering in General State-Space Models”, IEEE Conference on Decision and Control, Las Vegas 2002.

[9] V. B. Tadić and A. Doucet, “A Simulation Based Algorithm for Optimal Quantization of Hidden Markov Models,” International Symposium on Information Theory, Yokohama 2003.

[10] V. B. Tadić and A. Doucet, “A Simulation Based Algorithm for the Approximation of the Optimal Filters for Hidden Markov Models”, to be submitted.