RAO-BLACKWELLIZED PARTICLE FILTER FOR MARKOV MODULATED NONLINEARDYNAMIC SYSTEMS

Saikat Saha and Gustaf Hendeby

Saikat Saha and Gustaf Hendeby, RAO-BLACKWELLIZED PARTICLE FILTER FOR MARKOV MODULATED NONLINEARDYNAMIC SYSTEMS, 2014, IEEE statistical signal processing workshop (SSP14), Gold Coast, Australia.

RAO-BLACKWELLIZED PARTICLE FILTER FOR MARKOV MODULATED NONLINEAR DYNAMIC SYSTEMS

Saikat Saha[†,*]

[†] Dept. Electrical Engineering
Linköping University
SE-581 83 Linköping, Sweden.
{saha, hendeby}@isy.liu.se

Gustaf Hendeby[†,‡]

[‡] Dept. of Sensor & EW Systems
Swedish Defence Research Agency (FOI)
SE-581 11 Linköping, Sweden
gustaf.hendeby@foi.se

ABSTRACT

The Markov modulated (switching) state space is an important model paradigm in statistical signal processing. In this article, we specifically consider Markov modulated nonlinear state-space models and address the online Bayesian inference problem for such models. In particular, we propose a new Rao-Blackwellized particle filter for the inference task which is our main contribution here. A detailed description of the problem and an algorithm is presented.

Index Terms— Rao-Blackwellized particle filter, Markov regime switching, switching nonlinear state space, Jump Markov nonlinear systems

1. INTRODUCTION

Many practical applications in applied science often deals with nonlinear dynamic systems involving both a continuous value target state and a discrete value regime variable. Such descriptions imply that the system can switch between different nonlinear dynamic regimes, where the parameters of each regime is governed by the corresponding regime variable. The different regimes can possibly be described in terms of different stochastic processes. The regime variable also evolves dynamically according to a finite state (assumed) Markov chain. Both the target state and regime variable are latent and are related to noisy observations. This model paradigm is often referred to as a *Markov regime switching* (MRS) state space, sometimes with other monikers like jump Markov, Markov modulated, or hybrid dynamic system. Due to its modeling flexibility, MRS is very popular in different disciplines and as such, has been successfully used in diverse areas like econometrics, operations research, signal processing and machine learning among others [1–3]. However, most of the studies have focused on a special case where each individual regime follows a linear Gaussian state-space model. This special case is known as the *jump Markov linear system* (JMLS). Nonetheless, for many practical applications of interest, including econometrics [4], signal processing [5], target tracking and localization [6, 7], the individual regimes follows nonlinear dynamics, possibly driven by non-Gaussian processes. Such a system is referred to as a *Markov modulated nonlinear dynamic system* (MmNDS) or a *jump Markov nonlinear system* (JMNS). Compared to JMLS, this class of problems is less well studied. Hence, here we consider the state inference problem for MmNDS.

For certain models, part of the state space is (conditionally) tractable. It is then sufficient to employ a *particle filter* (PF) for the remaining intractable part of the state space. By exploiting such analytical substructure, the Monte Carlo based estimation is then confined to a space of lower dimension. Consequently, the estimate obtained is often better and never worse than the estimate provided by the PF targeting the full state space. This efficiency as a result of the well-known *Rao-Blackwell* estimator. For this reason, the method is popularly known as *Rao-Blackwellized particle filtering* (RBPF) [8].

In this article, based on [9], we address the online inference problem for MmNDS using PF. Particularly, we propose a new RBPF framework using the conditionally analytical substructure of the regime indicator variable. It is noted that a similar solution has recently been employed in the context of an expectation-maximization algorithm [10]; our paper differs in its emphasis on the RBPF framework as an independent method.

2. PROBLEM STATEMENT

2.1. Model description

Consider a (hybrid) nonlinear state-space model evolving according to

$$\Pi(r_k|r_{k-1}), \quad p_{\theta_{r_k}}(x_k|x_{k-1}, r_k), \quad p_{\theta_{r_k}}(y_k|x_k, r_k), \quad (1)$$

where $r_k \in S \triangleq \{1, 2, \dots, s\}$, is a (discrete) regime indicator variable with finite number of regimes (*i.e.*, categorical variable), $x_k \in \mathbb{R}^{n_x}$ is the (continuous) state variable. As the system can switch between different dynamic regimes,

*The authors would like to thank COOP-LOC, funded by SSF and CADICS, funded by Swedish Research Council (VR) for the financial supports.

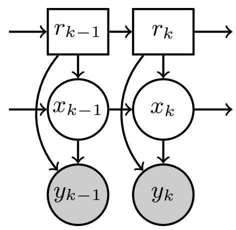

Fig. 1. Graphical representation of a Markov modulated nonlinear dynamic systems.

for a given regime variable $l \in S$, the corresponding dynamic regime can be characterized by a set of parameters θ_l. Both x_k and r_k are latent variables, related to the measurement $y_k \in \mathbb{R}^{n_y}$. The time behavior of the regime variable r_k is commonly modeled as a homogeneous (time-invariant) first order Markov chain with a *transition probability matrix* (TPM) $\Pi = [\pi_{ij}]_{ij}$ such that

$$\pi_{ij} \triangleq \mathbb{P}(r_k = j | r_{k-1} = i) \quad (i, j \in S), \tag{2}$$

where $\pi_{ij} \geq 0$ and $\sum_{j=1}^{s} \pi_{ij} = 1$. This model is represented graphically in Fig. 1. The following examples illustrate some real life applications where the above model is used.

Example 1: Consider the Markov switching stochastic volatility model [4], where x_k is the latent time varying log-volatility, y_k is the observed value of daily return of stock price or index. The regime variable r_k is modeled as a K*-state* first order Markov process. The model is further specified as

$$p_{\theta_{r_k}}(x_k | x_{k-1}, r_k) = \mathcal{N}(\alpha_{r_k} + \phi x_{k-1}, \sigma^2), \tag{3a}$$
$$p(y_k | x_k, r_k) = \mathcal{N}(0, e^{x_k/2}), \tag{3b}$$

where the parameter vector is given by $\theta_{r_k} \triangleq \{\alpha_{r_k}, \phi, \sigma\}$.

Example 2: Consider an altitude based terrain navigation framework [11]. To keep the description simple, assume that an aircraft is traveling in an one dimensional space (*e.g.*, on a manifold). The aircraft is assumed to follow a constant velocity model. The state-space model is given by

$$x_{k+1} = \begin{pmatrix} 1 & T \\ 0 & 1 \end{pmatrix} x_k + \begin{pmatrix} \frac{1}{2}T^2 \\ T \end{pmatrix} w_k \tag{4a}$$
$$y_k = h(x_k) + e_k(r_k), \tag{4b}$$

where T is the sampling period, w_k and $e_k(\cdot)$ are the process and the measurement noise, respectively, commonly assumed to be individually mutually independent. The aircraft's latent state x_k consists of position and velocity. The observation y_k denotes the terrain altitude measured by the aircraft at time k. This is obtained by deducting the height measurement of the ground looking (on board) radar from the known altitude of the aircraft (obtained using an altimeter). The function $h(x_k)$ relates the terrain altitude to position x_k in the form of a digital terrain database. $h(\cdot)$ is highly nonlinear (the measured height can corresponds to different locations) .

The distribution of w_k is typically modeled as Gaussian. As radar reflections can come from the ground as well as from the tree canopy, typically the observation noise e_k is modeled as a (bimodal) two component Gaussian mixture. The regime variable r_k indicates the corresponding mixture component. The sufficient statistics (*i.e.*, mean and variance) of each component can be specified by the regime dependent parameters θ_{r_k}. The dynamics of r_k is modeled as two state first order homogeneous Markov process.

2.2. Inference objective

For the model described by (1)–(2), given the densities for the initial state $\{r_0, x_0\}$ and the measurements up to time k ($Y_k \triangleq (y_1, \cdots, y_k)$), our interest lies in estimating sequentially the latent states $\{r_k, x_k\}$. More precisely, for the statistical inference purpose, we target the series of filtering distributions $\mathbb{P}(r_k | Y_k)$ and $p(x_k | Y_k)$ recursively over time. However, the above posteriors are in general, computationally intractable. Given this intractability, PF is a suitable candidate for this approximate (real time) inference task. We note that conditioned on the sequence $X_k \triangleq x_{1:k}$, r_k follows a finite state *hidden Markov model* (HMM), implying that $\mathbb{P}(r_k | X_k, Y_k)$ is analytically tractable. Using this analytical substructure, it is possible to implement an efficient RBPF scheme which can reduce the variance of the estimation error. In the sequel, we detail this RBPF framework for the MmNDS.

3. A NEW RBPF FOR MARKOV MODULATED NONLINEAR STATE-SPACE MODEL

3.1. Description of the RBPF approach

The initial densities for the state and the regime variables are given by $p(x_0)$ and $\mathbb{P}(r_0) \triangleq \mathbb{P}(r_0 | x_0)$, respectively, these can be arbitrary but are assumed to be known. We further assume favorable mixing conditions as in [12].

Suppose that we are at time $k-1$. We consider the extended target density $p(r_{k-1}, X_{k-1} | Y_{k-1})$ which can be decomposed as

$$p(r_{k-1}, X_{k-1} | Y_{k-1}) = p(r_{k-1} | X_{k-1}, Y_{k-1})\, p(X_{k-1} | Y_{k-1}). \tag{5}$$

The posterior propagation of the latent state x_{k-1} can then be targeted through a PF, where $p(X_{k-1} | Y_{k-1})$ is represented by a set of N weighted random particles as

$$p(X_{k-1} | Y_{k-1}) \approx \sum_{i=1}^{N} w_{k-1}^{(i)} \delta(X_{k-1} - X_{k-1}^{(i)}). \tag{6}$$

Conditioned on $\{X_{k-1}, Y_{k-1}\}$, the regime variable r_{k-1} follows a finite state-space HMM, making $p(r_{k-1} | X_{k-1}, Y_{k-1})$

analytically tractable[1], which is represented as

$$q_{k-1|k-1}^{(i)}(l) \triangleq \mathbb{P}(r_{k-1} = l | X_{k-1}^{(i)}, Y_{k-1}), \quad (7)$$

for $l \in S$ and $i = 1, \ldots, N$. Now using (6) and (7), the extended target density in (5) can be represented as

$$\left[X_{k-1}^{(i)}, w_{k-1}^{(i)}, \{q_{k-1|k-1}^{(i)}(l)\}_{l=1}^{s}\right]_{i=1}^{N}. \quad (8)$$

Now having observed y_k, we want to propagate the extended target density in (5) to time k. This can be achieved in the following steps (a)–(d):

(a) Prediction step for conditional HMM filter: this is easily obtained as

$$q_{k|k-1}^{(i)}(l) \triangleq \mathbb{P}(r_k = l | X_{k-1}^{(i)}, Y_{k-1}) \quad (9a)$$

$$= \sum_{j=1}^{s} \pi_{jl}\, q_{k-1|k-1}^{(i)}(j), \quad (l, j) \in S. \quad (9b)$$

(b) Prediction step for particle filter: at this stage, generate N new samples $x_k^{(i)}$ from an appropriate proposal kernel as

$$x_k^{(i)} \sim \pi(x_k | X_{k-1}^{(i)}, Y_k). \quad (10)$$

Then set $X_k^{(i)} = \{X_{k-1}^{(i)}, x_k^{(i)}\}$, for $i = 1, \ldots, N$, representing the particle trajectories up to time k.

(c) Update step for conditional HMM filter: noting that

$$\mathbb{P}(r_k = l | X_k, Y_k) \propto$$
$$p(y_k, x_k | r_k = l, X_{k-1}, Y_{k-1})\, \mathbb{P}(r_k = l | X_{k-1}, Y_{k-1}), \quad (11)$$

we have

$$q_{k|k}^{(i)}(l) \propto p(y_k, x_k^{(i)} | r_k = l, X_{k-1}^{(i)}, Y_{k-1})\, q_{k|k-1}^{(i)}(l)$$
$$\propto p_{\theta_l}(y_k | x_k^{(i)}, r_k = l)\, p_{\theta_l}(x_k^{(i)} | x_{k-1}^{(i)}, r_k = l)\, q_{k|k-1}^{(i)}(l). \quad (12)$$

Now defining

$$\alpha_k^{(i)}(l) \triangleq p_{\theta_l}(y_k | x_k^{(i)}, r_k = l)\, p_{\theta_l}(x_k^{(i)} | x_{k-1}^{(i)}, r_k = l)\, q_{k|k-1}^{(i)}(l) \quad (13)$$

we obtain

$$q_{k|k}^{(i)}(l) = \alpha_k^{(i)}(l) \Big/ \textstyle\sum_{j=1}^{s} \alpha_k^{(i)}(j), \quad (14)$$

for $l \in S$ and $i = 1, \ldots, N$.

(d) Update step for particle filter: as the continuous state can be recursively propagated using the following relation:

$$p(X_k | Y_k) \propto p(y_k, x_k | X_{k-1}, Y_{k-1})\, p(X_{k-1} | Y_{k-1}), \quad (15)$$

the corresponding weight update equation for the particle filtering is given by

$$w_k^{(i)} = \frac{p(x_k^{(i)}, y_k | X_{k-1}^{(i)}, Y_{k-1})}{\pi_k(x_k^{(i)} | X_{k-1}^{(i)}, Y_k)} \widetilde{w}_{k-1}^{(i)} \quad (16a)$$

$$\widetilde{w}_k^{(i)} = w_k^{(i)} \Big/ \textstyle\sum_{j=1}^{N} w_k^{(j)}, \quad (16b)$$

where $\{\widetilde{w}_k^{(i)}\}_{i=1}^{N}$ are the normalized weights. The numerator $p(x_k^{(i)}, y_k | X_{k-1}^{(i)}, Y_{k-1})$ can be obtained as

$$p(x_k^{(i)}, y_k | X_{k-1}^{(i)}, Y_{k-1}) = \sum_{l=1}^{s} p(x_k^{(i)}, y_k | r_k = l, X_{k-1}^{(i)}, Y_{k-1}) \times$$
$$\mathbb{P}(r_k = l | X_{k-1}^{(i)}, Y_{k-1}), \quad (17)$$

which is basically given by the normalizing constant of (14). Note that the marginal density $p(x_k | Y_k)$ can be obtained as

$$p(x_k | Y_k) \approx \sum_{i=1}^{N} \widetilde{w}_k^{(i)} \delta(x_k - x_k^{(i)}). \quad (18)$$

The posterior probability of the regime variable can now be obtained as

$$\mathbb{P}(r_k = l | Y_k) \approx \sum_{i=1}^{N} q_{k|k}^{(i)}(l) \widetilde{w}_k^{(i)}. \quad (19)$$

Let $\hat{m}_k^{(i)}$ and $\hat{V}_k^{(i)}$ denote the mean and variance of the conditional HMM filter. They are now obtained as

$$\hat{m}_k^{(i)} = \sum_{l=1}^{s} (r_k = l)\, q_{k|k}^{(i)}(l), \quad (20a)$$

$$\hat{V}_k^{(i)} = \sum_{l=1}^{s} \big((r_k = l) - \hat{m}_k^{(i)}\big)(\cdot)^T q_{k|k}^{(i)}(l), \quad (20b)$$

where $(A)(\cdot)^T$ is shorthand for AA^T. As noted earlier, the posterior of the regime variable is given by (19). Let $\hat{m}_k$ and $\hat{V}_k$ denote the corresponding mean and variance, which can be obtained as

$$\hat{m}_k = \sum_{i=1}^{N} \widetilde{w}_k^{(i)}\, \hat{m}_k^{(i)}, \quad (21a)$$

$$\hat{V}_k = \sum_{i=1}^{N} \widetilde{w}_k^{(i)} \Big(\hat{V}_k^{(i)} + (\hat{m}_k^{(i)} - \hat{m}_k)(\cdot)^T\Big). \quad (21b)$$

Remark 1 *For PF, a popular (but less efficient) choice for the proposal kernel is given by the state transition density $p(x_k | x_{k-1})$, which in this case can be obtained in the form of a weighted mixture density:*

$$p(x_k | x_{k-1}^{(i)}) = \sum_{l=1}^{s} p_{\theta_{r_k}}(x_k | x_{k-1}^{(i)}, r_k = l)\, q_{k|k-1}^{(i)}(l), \quad (22)$$

where $p_{\theta_{r_k}}(x_k | x_{k-1}^{(i)}, r_k = l)$ is specified in (1).

The new RBPF for the MmNDS is summarized in Alg. 1

4. RELATION TO OTHER SIMILAR MODELS

Conditionally finite state-space HMM similar to the one presented here have previously been considered by [13] and [14],

[1] The 'favorable' mixing property ensures that $p(x_{0:k-1} | y_{1:k-1})$ can be well approximated by $p(x_{k-L:k-1} | y_{1:k-1})$, for some lag L. Consequently, $p(r_{k-1} | x_{0:k-1}, y_{1:k-1}) \approx p(r_{k-1} | x_{k-L:k-1}, y_{1:k-1})$.

Alg. 1 RBPF for MmNDM

Initialization:

For each particle $i = 1, \ldots, N$ do

- Sample $x_0^{(i)} \sim p(x_0)$,
- Set initial weights $w_0^{(i)} = \frac{1}{N}$,
- Set initial $q_{0|0}^{(i)}(l) \triangleq \mathbb{P}(r_0 = l | x_0^{(i)}), \quad l = 1, \ldots, s$

Iterations:

Set the resampling threshold η;

For $k = 1, 2, \ldots$ do

- For each particle $i = 1, \ldots, N$ do
  - Compute $q_{k|k-1}^{(i)}(l)$ using (9b)
  - Sample $x_k^{(i)} \sim \pi(x_{k|.})$ using (10)
  - Set $X_k^{(i)} \triangleq (X_{k-1}^{(i)}, x_k^{(i)})$
  - Compute $\alpha_k^{(i)}(l)$ using (13)
  - Compute $q_{k|k}^{(i)}(l)$ using (14)
  - Compute $w_k^{(i)}$ using (16a) and (17) as
  $$w_k^{(i)} = \sum_{j=1}^{s} \alpha_k^{(i)}(j) \Big/ \pi_k(x_k^{(i)} | X_{k-1}^{(i)}, Y_k) \cdot \widetilde{w}_{k-1}^{(i)}$$
- Normalize the weights using (16b)
- Compute $N_{\text{eff}} = 1 / \sum_{i=1}^{N} (\widetilde{w}_k^{(i)})^2$.
  - If $N_{\text{eff}} \leq \eta$, resample the particles. Let the resampled particles be $i^* = 1, \ldots, N$.
  - Copy the corresponding $q_{k|k}^{(i^*)}(l)$ and set $\widetilde{w}_k^{(i^*)} = \frac{1}{N}$.

although, each framework is fundamentally different. The differences are emphasized below.

In the model in this paper, (x_k, y_k) follows a nonlinear state-space model, modulated by a finite state *hidden Markov process* (HMP) r_k. Hierarchically r_k is at the top level and is not influenced by x_k. This is different from the hierarchical conditionally finite state-space HMM in [13], where (r_k, y_k) follows a finite state-space HMP, which is modulated by a another (hidden) Markov process c_k. Here c_k is at the top of hierarchy and is not influenced by r_k. In contrast, [14] considered a partially observable finite state-space HMM, where r_k is a finite state HMP, y_k is a latent data process and z_k is observed data process. Conditioned on the sequence $z_{1:k}$, here (r_k, y_k) follows a finite state-space HMM.

5. CONCLUDING REMARKS

Markov modulated nonlinear state-space model, although less well explored, appears naturally in many applications of interest. The model implies that the system can switch between different nonlinear dynamic regimes. The regime state is governed by a regime variable, which follows a homogeneous finite state first-order Markov process. Here, the associated online inference problem for this model is addressed. In particular, a new RBPF is proposed for the inference. This RBPF scheme exploits analytical marginalization of the regime variable using the conditional HMM structure. This results in improved performance over a standard particle filter in terms of variance of the estimation error. Moreover for a standard particle filter where the regime state is also represented by the particles, degeneracy is commonly observed around regime transition [6]. In our RBPF implementation, as the regime variable follows a conditionally analytical substructure, the degeneracy is expected to be less severe.

References

[1] S. Fruhwirth-Schnatter, *Finite Mixture and Markov Switching Models*, Springer, 2007.

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

[3] D. Barber, *Bayesian Reasoning and Machine Learning*, Cambridge University Press, 2012.

[4] C. M. Carvalho and H. F. Lopes, "Simulation-based sequential analysis of Markov switching stochastic volatility models," *Computational Statistics & Data Analysis*, vol. 51, pp. 4526–4542, July 2003.

[5] C. Andrieu, M. Davy, and A. Doucet, "Efficient particle filtering for jump Markov systems," *IEEE Transactions on Signal Processing*, vol. 51, no. 7, pp. 1762–1770, July 2003.

[6] H. Driessen and Y. Boers, "Efficient particle filter for jump Markov nonlinear systems," *IEE Proceedings- radar, sonar and navigation*, vol. 152, no. 5, pp. 323–326, 2005.

[7] S. A. Pasha, B. N. Vo, H. D. Tuan, and W. K. Ma, "A Gaussian mixture PHD filter for jump Markov system models," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 45, no. 3, pp. 919–936, 2009.

[8] O. Cappé, S. J. Godsill, and E. Moulines, "An overview of existing methods and recent advances in sequential Monte Carlo," *Proceedings of IEEE*, vol. 95 (5), pp. 899–924, 2007.

[9] S. Saha and G. Hendeby, "Online inference in Markov modulated nonlinear dynamic systems: a Rao-Blackwellized particle filtering approach," *ArXiv: 1311.6486v1*, Nov. 2013.

[10] E. Özkan, F. Lindsten, C. Fritsche, and F. Gustafsson, "Recursive maximum likelihood identification of jump markov nonlinear systems," *ArXiv: 1312.0781v1*, Dec. 2013.

[11] T. Schön, F. Gustafsson, and P. J. Nordlund, "Marginalized particle filter for mixed linear/nonlinear state space models," *IEEE Transaction on Signal Processing*, vol. 53, no. 7, pp. 2279–2289, 2005.

[12] D. Crisan and A. Doucet, "A survey of convergence results on particle filtering methods for practitioners," *IEEE Transaction on Signal Processing*, vol. 50, no. 3, pp. 736–746, 2002.

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

[14] C. Andrieu and A. Doucet, "Particle filtering for partially observed Gaussian state space models," *Journal of the Royal Statistical Society Series B*, vol. 64, no. 4, pp. 827–836, 2002.