



HAL
open science

Backward importance sampling for partially observed diffusion processes

Marie-Pierre Etienne, Pierre Gloaguen, Sylvain Le Corff, Jimmy Olsson

► **To cite this version:**

Marie-Pierre Etienne, Pierre Gloaguen, Sylvain Le Corff, Jimmy Olsson. Backward importance sampling for partially observed diffusion processes. 2020. hal-02476102v1

HAL Id: hal-02476102

<https://hal.science/hal-02476102v1>

Preprint submitted on 12 Feb 2020 (v1), last revised 7 May 2021 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Backward importance sampling for partially observed diffusion processes

Marie-Pierre Étienne^{*}, Pierre Gloaguen^{*}, Sylvain Le Corff[†], and Jimmy Olsson[‡]

^{*}Institut de recherche mathématique de Rennes, IRMAR, UMR 6625.

^{*}AgroParisTech, UMR MIA 518.

[†]Samovar, Télécom SudParis, Département CITI, TIPIC, Institut Polytechnique de Paris.

[‡]Department of Mathematics, KTH Royal Institute of Technology, Stockholm.

Abstract

This paper proposes a new Sequential Monte Carlo algorithm to perform maximum likelihood estimation in partially observed diffusion processes. Training such generative models and obtaining low variance estimators of the posterior distributions of the latent states given the observations is challenging as the transition densities of the latent states cannot be evaluated pointwise. In this paper, a backward importance sampling step is introduced to estimate such posterior distributions instead of the usual acceptance-rejection approach. This allows to use unbiased estimates of the unknown transition densities available under mild assumptions for multivariate stochastic differential equations while acceptance-rejection based methods require strong conditions to obtain upper-bounded estimators. The performance of this estimator is assessed in the case of a partially observed stochastic Lotka-Volterra model.

1 Introduction

Latent data models are all-pervasive in time series and sequential data analysis across a wide range of applied science and engineering domains such as movement ecology [Michelot et al., 2016], energy consumptions modelling [Candanedo et al., 2017], genomics [Yau et al., 2011, Gassiat et al., 2016, Wang et al., 2017], target tracking [Särkkä et al., 2007], enhancement and segmentation of speech and audio signals [Rabiner, 1989], see also [Särkkä, 2013, Douc et al., 2014, Zucchini et al., 2017] and the numerous references therein. The data considered in this paper originate from partially observed diffusion (POD) processes. The observations are assumed to be defined as random functions of a continuous-time diffusion process. This diffusion process $(X_t)_{t \geq 0}$ is the solution to a Stochastic Differential Equation (SDE) driven by a Brownian motion so that any discrete sub-sample $(X_k)_{0 \leq k \leq n}$ is a Markov chain. From this perspective, a POD is a general Hidden Markov Model (HMM).

In this setting, performing maximum likelihood estimation (MLE) for instance with the Expectation Maximization (EM) algorithm [Dempster et al., 1977] or a stochastic gradient ascent ([Cappé et al., 2005] in the case of HMMs) is a challenging task. Both approaches involve conditional distributions of sequences of hidden states given the observation record (the *smoothing* distribution), which are not available explicitly. Markov chain Monte Carlo (MCMC) and sequential Monte Carlo (SMC) methods (also known as particle

filters or smoothers) are widespread solutions to propose consistent estimators of such distributions. However, a pivotal tool of both MCMC and SMC approaches is the evaluation of the density of the conditional distribution of X_k given X_{k-1} for all $1 \leq k \leq n$. Except in very few SDE cases, this transition density has no analytical expression and the key conditional expectations cannot be computed explicitly so that MLE cannot be achieved. The objective of this paper is to propose an efficient consistent and asymptotically Gaussian SMC algorithm which can be applied to general POD processes to approximate conditional expectations of functionals of the hidden states.

Following [Fearnhead et al., 2008] and [Gloaguen et al., 2018, Gloaguen et al., 2019], this paper concentrates on SMC methods to approximate smoothing distributions with a random set of states, the *particles*, associated with importance weights by combining importance sampling and resampling steps. The online algorithm of [Gloaguen et al., 2018] may be used to approximate expectations of additive functionals under the smoothing distributions by processing the data stream online. This algorithm extends the particle-based rapid incremental smoother (PaRIS) of [Olsson et al., 2017] when the unknown transition densities are replaced by unbiased estimates. This approach is an online version of the Forward Filtering Backward Simulation algorithm [Douc et al., 2011] specifically designed to approximate conditional expectations of additive functionals. The crucial feature which makes the PaRIS algorithm appealing is the acceptance-rejection step which benefits from the unbiased estimation. The extension of the usual alternative, named the Forward Filtering Backward Smoothing algorithm [Doucet et al., 2000], is more sensitive as it involves ratios of these unknown quantities. Other smoothing algorithms such as two-filter based approaches [Briers et al., 2010, Fearnhead et al., 2010b, Nguyen et al., 2017] could be extended similarly but they are intrinsically not online procedures as they require the time horizon and all observations to be available to initialize a backward information filter. In [Gloaguen et al., 2018, Gloaguen et al., 2019], the consistency of this algorithm as long as a central limit theorem (CLT) are established. This makes this pseudo marginal smoother the first algorithm to approximate such expectations in the general setting of this paper with theoretical guarantees and an explicit expression of the asymptotic variance. This work extended the result of [Olsson et al., 2017], written only in the case of the bootstrap filter [Gordon et al., 1993], and the theoretical guarantees obtained for online sequential Monte Carlo smoothers given in [Del Moral et al., 2010, Douc et al., 2011, Dubarry and Le Corff, 2013, Gerber and Chopin, 2017].

However, a requirement to use this algorithm is that the unbiased estimate of the transition is almost surely positive and upper bounded. This condition is required to perform a pivotal backward acceptance-rejection sampling procedure. In the context of diffusion processes, this assumption is very restrictive and narrows the possible models to the class of diffusions satisfying the Exact algorithm conditions of [Beskos et al., 2006a], for which General Poisson Estimators (GPEs) [Fearnhead et al., 2008] lead to eligible unbiased estimators.

In this paper, a new procedure is introduced to replace the backward acceptance-rejection step by an importance sampling estimate which leads to a smoothing algorithm that only requires an almost surely positive estimator of the transition density. Moreover, it is shown that such an estimator can be obtained for a wide range of diffusion processes, using the parametrix estimators of [Andersson et al., 2017] and [Fearnhead et al., 2017]. The overall method therefore provides a consistent estimator of expectations under the smoothing distributions in the context of generic POD, where the hidden diffusion process can belong to a wide range of multivariate models. Moreover, this method does not rely on any acceptance rejection procedure which reduces very significantly the computational time and its variance, as illustrated in the numerical experiments.

2 Model and objectives

Let $\Theta \subset \mathbb{R}^q$ be a compact parameter space and $(X_t)_{t \geq 0}$ be defined as a weak solution to the following SDE in \mathbb{R}^d :

$$dX_t = \alpha_\theta(X_t)dt + \sigma_\theta(X_t)dW_t, \quad (1)$$

where $\theta \in \Theta$, $(W_t)_{t \geq 0}$ is a standard Brownian motion in \mathbb{R}^d , $\alpha_{\theta_*} : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is the drift function and $\sigma_{\theta_*} : \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}$ is the diffusion. It is assumed that the solution to (1) is partially observed at times $t_0 = 0, \dots, t_n$, for a given $n \geq 1$, through an observation process $(Y_k)_{0 \leq k \leq n}$ taking values in \mathbb{R}^m . For all $0 \leq k \leq n$, the distribution of Y_k given $(X_t)_{t \geq 0}$ depends on $X_k = X_{t_k}$ only and has density $g_{k;\theta_*}$ with respect to the Lebesgue measure. The distribution of X_0 has density χ with respect to the Lebesgue measure and for all $0 \leq k \leq n-1$, the conditional distribution of X_{k+1} given $(X_t)_{0 \leq t \leq t_k}$ has density $q_{k+1;\theta_*}(X_k, \cdot)$.

In this setting, common learning objectives are the *state estimation problem*, which aims at recovering the underlying signal X_k at time t_k given the observations $Y_{0:n}$, where $a_{u:v}$ is a short-hand notation for (a_u, \dots, a_v) , and the *parameter inference problem* which aims at approximating

$$\hat{\theta}_n = \operatorname{argmax}_{\theta \in \Theta} L_n(\theta),$$

where $L_n(\theta)$ is the likelihood of the observations. When θ is known, the state estimation problem is usually solved by approximating the posterior mean of X_k given the observations $Y_{0:n}$ when the model is driven by the parameter θ . In the context of parameter estimation, note that

$$L_n(\theta) = \int \chi(x_0) g_{0;\theta}(x_0, Y_0) \prod_{k=0}^{n-1} r_{k;\theta}(x_k, x_{k+1}) dx_{0:n},$$

where, for all $0 \leq k \leq n$ and all $\theta \in \Theta$,

$$r_{k;\theta}(x_k, x_{k+1}) = q_{k+1;\theta}(x_k, x_{k+1}) g_{k+1;\theta}(x_{k+1}, Y_{k+1}).$$

Expectation Maximization based algorithms are appealing solutions to obtain an estimator of $\hat{\theta}_n$. The pivotal concept of the EM algorithm is that the intermediate quantity defined by

$$\theta \mapsto Q(\theta, \theta') = \mathbb{E}_{\theta'} \left[\sum_{k=0}^{n-1} \log r_{k;\theta}(X_k, X_{k+1}) \middle| Y_{0:n} \right] \quad (2)$$

may be used as a surrogate for $L_n(\theta)$ in the maximization procedure, where $\mathbb{E}_{\theta'}$ is the expectation under the joint distribution of the latent states and the observations when the model is parameterized by θ . Therefore, the EM algorithm iteratively builds a sequence $(\theta_p)_{p \geq 0}$ of parameter estimates following the two steps: (i) compute $\theta \mapsto Q(\theta, \theta_p)$ and (ii) choose θ_{p+1} as a maximizer of $\theta \mapsto Q(\theta, \theta_p)$. Gradient ascent algorithms are compelling alternatives to the EM algorithm. In the context of HMMs, the gradient of the log-likelihood can also be expressed as an expectation of an additive functional of the hidden states given $Y_{0:n}$ ([Cappé et al., 2005], Chapter 10, or [Gloaguen et al., 2019] in the context of SDEs). A key feature here is that all the relevant estimators rely on computing, for some parameters θ and θ' :

$$\mathbb{E}_{\theta'} [h_{0:n,\theta}(X_{0:n}) | Y_{0:n}],$$

where $h_{0:n,\theta}$ is an *additive functional*, i.e. satisfying:

$$h_{0:n,\theta} : x_{0:n} \mapsto \sum_{k=0}^{n-1} \tilde{h}_{k;\theta}(x_k, x_{k+1}),$$

where $\tilde{h}_{k;\theta}$ is a functional depending on the estimator. For any $\theta \in \Theta$, $0 \leq k_1 \leq k_2 \leq n$ and any bounded and measurable function h on $(\mathbb{R}^d)^{k_2-k_1+1}$, define the *joint smoothing distributions* as:

$$\phi_{k_1:k_2|n;\theta}[h] := \mathbb{E}_\theta [h(X_{k_1:k_2})|Y_{0:n}] . \quad (3)$$

For all $0 \leq k \leq n$, $\phi_{k;\theta} = \phi_{k:k|k;\theta}$ are the *filtering distributions*.

In the following, θ is dropped from the notations for better clarity when there is no possible confusion. As noted for instance in [Cappé et al., 2005], although the objective is to obtain approximation of smoothing distributions, the filtering distribution is crucial as, for additive functionals,

$$\phi_{0:n|n}[h_{0:n}] = \phi_n [\mathbf{T}_n[h_{0:n}]] ,$$

where

$$\mathbf{T}_n[h_{0:n}](X_n) = \mathbb{E} [h_n(X_{0:n})|X_n, Y_{0:n}] . \quad (4)$$

As a key consequence of the additive property, for all $1 \leq k \leq n$

$$\mathbf{T}_k[h_{0:k}](X_k) = \mathbb{E} \left[\mathbf{T}_{k-1}[h_{0:(k-1)}](X_{k-1}) + \tilde{h}_{k-1}(X_{k-1}, X_k) \middle| X_k, Y_{0:k-1} \right] . \quad (5)$$

However, the exact computation of all these key expectations is not possible in general state spaces. The next section describes a Sequential Monte Carlo algorithm [Doucet et al., 2013] to approximate ϕ_n by weighted samples $\{(\omega_n^\ell, \xi_n^\ell)\}_{\ell=1}^N$ and the algorithm of [Gloaguen et al., 2018] to compute recursively, for each $1 \leq \ell \leq N$ an approximation τ_n^ℓ of $\mathbf{T}_n[h_{0:n}](\xi_n^\ell)$ so that the estimator of $\phi_{0:n|n}[h_{0:n}]$ is defined as

$$\phi_{0:n|n}^N[h_{0:n}] := \sum_{\ell=1}^N \frac{\omega_n^\ell}{\sum_{j=1}^N \omega_n^j} \tau_n^\ell . \quad (6)$$

3 Online sequential Monte Carlo smoothing

In the case of POD processes, SMC methods cannot be used straightforwardly as the transition densities q_k , $0 \leq k \leq n-1$, are unknown. To overcome these issues, following [Fearnhead et al., 2008, Olsson et al., 2011, Gloaguen et al., 2018], consider the following assumption. Let $(\mathcal{U}, \mathcal{B}(\mathcal{U}))$ be a general state space.

H1 For all $\theta \in \Theta$ and $k \geq 0$, there exists a Markov kernel on $(\mathbb{R}^d \times \mathbb{R}^d, \mathcal{B}(\mathcal{U}))$ with density $\mathbf{K}_{k;\theta}$ with respect to a reference measure μ on $(\mathcal{U}, \mathcal{B}(\mathcal{U}))$ and a positive mapping $\bar{r}_{k;\theta}$ on $\mathbb{R}^d \times \mathbb{R}^d \times \mathcal{U}$ such that, for all $(x, x') \in \mathbb{R}^d \times \mathbb{R}^d$,

$$\int \mathbf{K}_{k;\theta}(x, x'; z) \bar{r}_{k;\theta}(x, x'; z) \mu(dz) = r_{k;\theta}(x, x') .$$

3.1 Filtering

Let $(\xi_0^\ell)_{\ell=1}^N$ be independent and identically distributed according to an instrumental proposal density ρ_0 on \mathbb{R}^d and define the importance weights $\omega_0^\ell := \chi(\xi_0^\ell)/\rho_0(\xi_0^\ell)$, where χ is the density of the distribution of X_0 , see Section 2. For any bounded and measurable function f defined on \mathbb{R}^d ,

$$\phi_0^N[f] := \Omega_0^{-1} \sum_{\ell=1}^N \omega_0^\ell f(\xi_0^\ell) , \quad \text{where} \quad \Omega_0 := \sum_{\ell=1}^N \omega_0^\ell .$$

is a consistent estimator of $\phi_0[f]$. Then, for all $k \geq 1$, once the observation Y_k is available, the weighted particle sample $\{(\omega_{k-1}^\ell, \xi_{k-1}^\ell)\}_{\ell=1}^N$ is transformed into a new weighted particle sample approximating ϕ_k . This update step is carried through in two steps, *selection* and *mutation*, using the auxiliary sampler introduced in [Pitt and Shephard, 1999]. New indices and particles $\{(I_k^\ell, \xi_k^\ell, \zeta_k^\ell)\}_{\ell=1}^N$ are simulated independently from the instrumental distribution with density on $\{1, \dots, N\} \times \mathbb{R}^d \times \mathcal{U}$:

$$\nu_k(\ell, x, z) \propto \omega_{k-1}^\ell \vartheta_{k-1}(\xi_{k-1}^\ell) p_{k-1}(\xi_{k-1}^\ell, x) \mathbf{K}_k(\xi_{k-1}^\ell, x; z),$$

where ϑ_{k-1} is an adjustment multiplier weight function and p_{k-1} a Markovian transition density. In practice, this step is performed as follows.

1. Sample I_k^ℓ in $\{1, \dots, N\}$ with probabilities proportional to $\{\omega_{k-1}^j \vartheta_{k-1}(\xi_{k-1}^j)\}_{1 \leq j \leq N}$.
2. Sample ξ_k^ℓ with distribution $p_{k-1}(\xi_{k-1}^{I_k^\ell}, \cdot)$.
3. Sample ζ_k^ℓ with distribution $\mathbf{K}_k(\xi_{k-1}^{I_k^\ell}, \xi_k^\ell; \cdot)$.

The choice of the proposal distribution p_{k-1} is a pivotal tuning step to obtain efficient estimations of the filtering distributions. In the context of this paper, a natural choice for p_k could be to use an approximation of the bootstrap filter, i.e. an approximation of q_{k-1} based for instance on a Euler discretization scheme. In the numerical section of this paper, we propose to use an approximation of the optimal filter which accounts for the newly obtained observation to propose new particles, see (12). For any $\ell \in \{1, \dots, N\}$, ξ_k^ℓ is associated with the importance weight defined by:

$$\omega_k^\ell := \frac{\bar{r}_{k-1}(\xi_{k-1}^\ell, \xi_k^\ell; \zeta_k^\ell)}{\vartheta_{k-1}(\xi_{k-1}^{I_k^\ell}) p_{k-1}(\xi_{k-1}^{I_k^\ell}, \xi_k^\ell)} \quad (7)$$

to produce the following approximation of $\phi_k[f]$:

$$\phi_k^N[f] := \Omega_k^{-1} \sum_{\ell=1}^N \omega_k^\ell f(\xi_k^\ell), \quad \text{where} \quad \Omega_k := \sum_{\ell=1}^N \omega_k^\ell.$$

3.2 Smoothing

In the context of additive functionals, the forward-only smoothing algorithm introduced in [Del Moral et al., 2010] proposes a particle approximation of (4) that can be computed *online* using the recursion (5). This algorithm has a computational complexity which grows *quadratically* with the number of particles N . This computational cost can be reduced when the transition density of the hidden states is upper bounded following [Olsson et al., 2017] by applying the accept-reject sampling approach proposed in [Douc et al., 2011] and illustrated in [Dubarry and Le Corff, 2011]. Following [Gloaguen et al., 2019], the backward statistics $\mathbf{T}_{k+1}[h_{0:k+1}](\xi_{k+1}^i)$, where \mathbf{T}_{k+1} is defined in (4), are estimated, for all $1 \leq i \leq N$, as follows,

$$\tau_{k+1}^i = \frac{1}{\tilde{N}} \sum_{j=1}^{\tilde{N}} \left(\tau_k^{J_{k+1}^{(i,j)}} + \tilde{h}_k \left(\xi_k^{J_{k+1}^{(i,j)}}, \xi_{k+1}^i \right) \right),$$

where $\tilde{N} \geq 1$ is a sample size which is typically small compared to N and where $(J_{k+1}^{(i,j)}, \xi_{k+1}^i)$, $1 \leq j \leq \tilde{N}$, are i.i.d. in $\{1, \dots, N\} \times \mathcal{U}$ with distribution

$$\bar{\nu}_k^i(\ell, z) \propto \omega_k^\ell \bar{r}_k(\xi_k^\ell, \xi_{k+1}^i; z) \mathbf{K}_k(\xi_k^\ell, \xi_{k+1}^i; z).$$

In [Gloaguen et al., 2018], it is assumed that, for all $0 \leq k \leq n$ and $0 \leq i \leq N$, there exists an upper bound $\bar{\varepsilon}_k^i$ such that

$$\sup_{\ell, \zeta} \bar{r}_k(\xi_k^\ell, \xi_{k+1}^i; \zeta) \leq \bar{\varepsilon}_k^i. \quad (8)$$

Then, for all $(i, z) \in \{1, \dots, N\} \times \mathbf{U}$,

$$\omega_k^\ell \bar{r}_k(\xi_k^\ell, \xi_{k+1}^i; z) \mathbf{K}_k(\xi_k^\ell, \xi_{k+1}^i; z) \leq \bar{\varepsilon}_k \omega_k^\ell \mathbf{K}_k(\xi_k^\ell, \xi_{k+1}^i; z).$$

Therefore, the following accept-reject mechanism algorithm may be used to sample from \bar{v}_k^i .

1. A candidate (J^*, ζ^*) is sampled in $\{1, \dots, N\} \times \mathbf{U}$ as follows:
 - (a) J^* is sampled with probabilities proportional to $(\omega_k^\ell)^N$;
 - (b) ζ^* is sampled independently with distribution $\mathbf{K}_k(\xi_k^{J^*}, \xi_{k+1}^i; \zeta^*)$.
2. (J^*, ζ^*) is then accepted with probability $\bar{r}_k(\xi_k^{J^*}, \xi_{k+1}^i; \zeta^*) / \bar{\varepsilon}_k$ and, upon acceptance,

$$J_{k+1}^{(i,j)} = J^*.$$

3.3 Unbiased estimators of the transition densities

The algorithm described above strongly relies on assumption H1. In the context of SDEs, when $g_{k+1; \theta}$ is available explicitly, this boils down to finding an unbiased estimate $\hat{q}_{k+1; \theta}(x, y; \zeta)$ of $q_{k+1; \theta}(x, y)$ and defining

$$\bar{r}_{k; \theta}(x, y; \zeta) = \hat{q}_{k+1; \theta}(x, y; \zeta) g_{k+1; \theta}(x_{k+1}, Y_{k+1}).$$

3.3.1 General Poisson Estimators

In [Olsson et al., 2011], [Gloaguen et al., 2018] and [Gloaguen et al., 2019], General Poisson Estimators (GPEs) are used to obtain an unbiased estimate of the transition density. However, designing such estimators requires three strong assumptions [Beskos et al., 2006a]:

1. the diffusion defined by (1) can be transformed into a unit diffusion through the Lamperti transform, with drift function $\tilde{\alpha}_\theta(x)$;
2. the drift of this unit diffusion can be expressed as the gradient of a potential function, i.e., there exists a twice differentiable function $A_\theta : \mathbb{R}^d \rightarrow \mathbb{R}$ such $\tilde{\alpha}_\theta = \nabla_x A_\theta$
3. The function $x \mapsto (\|\tilde{\alpha}_\theta(x)\|^2 + \Delta A_\theta(x)) / 2$ (where Δ denotes the Laplacian) is lower bounded.

Assumption (1) is used to define a proposal distribution absolutely continuous with respect to the target which is easy to sample from. Assumption (2) is necessary to obtain a tractable Radon-Nikodym derivative between the proposal and the target distributions using the Girsanov transformation. While these assumptions can be proved under mild assumptions for scalar diffusions, much stronger conditions are required in the multidimensional case [Aït-Shalia, 2008].

3.3.2 Parametrix estimators

More recently, [Andersson et al., 2017] and [Fearnhead et al., 2017] proposed an algorithm which can be used under weaker assumptions. This parametrix algorithm draws weighted skeletons using an importance sampling mechanism for diffusion processes. In this case, the sampled paths are not distributed as the target process but the weighted samples produce unbiased estimates of expectations of functionals of this process. To obtain an unbiased estimator $\hat{q}_{k+1}(x, y; \zeta)$, the parametrix algorithm draws weighted skeletons at random times $s_0 = 0 < s_1 < \dots < s_j$, denoted by $\{(x_{s_j}, w_{s_j})\}_{j \geq 0}$, where $x_0 = x$ and $w_0 = 1$. The update times $(s_j)_{j \geq 0}$ are instances of an inhomogeneous Poisson process of intensity $\lambda(t)$. Let (x_{s_j}, w_{s_j}) be the last weighted sample and s_{j+1} be the next update time of the trajectory. While $s_{j+1} < \Delta t_k$, the new state is sampled using a simple Euler scheme, namely:

$$x_{s_{j+1}} := x_{s_j} + \Delta s_j \alpha_\theta(x_{s_j}) + (\Delta s_j)^{1/2} \sigma_\theta(x_{s_j}) \varepsilon_{j+1},$$

where $\Delta s_j := s_{j+1} - s_j$, $\Delta t_k = t_{k+1} - t_k$ and $\varepsilon_{j+1} \sim \mathcal{N}_d(0, I_d)$. The proposal density associated with this procedure is denoted by $m_{j;\theta}(x_{s_j}, \cdot, \Delta \tau_j)$. Let \mathcal{K}^θ (resp. $\mathcal{K}_{\text{prop}}^{j,\theta}$) denote the Kolmogorov forward operator of the diffusion (resp. the Kolmogorov forward operator of the proposal distribution $m_{j;\theta}(x_{s_j}, \cdot, \Delta s_j)$). The forward operators write, for any function $h : \mathbb{R}^d \rightarrow \mathbb{R}$,

$$\mathcal{K}^\theta h(y) := - \sum_{i=1}^d \frac{\partial}{\partial y_i} \{ \alpha_{\theta,i}(y) h(y) \} + \sum_{i,\ell=1}^d \frac{1}{2} \frac{\partial^2}{\partial y_i \partial y_\ell} \{ \gamma_{\theta,i,\ell}(y) h(y) \},$$

where $\gamma_\theta = \sigma_\theta \sigma_\theta^T$. Then, following [Fearnhead et al., 2017], the weight is updated by

$$w_{s_{j+1}} := w_{s_j} \rho_j^\lambda(x_{s_j}, x_{s_{j+1}}, \Delta s_j),$$

where

$$\rho_j^\lambda(x, y, u) := 1 + \frac{(\mathcal{K} - \mathcal{K}_{\text{prop}}^{j,\theta}) m_{j;\theta}(x, z, u)|_{z=y}}{\lambda(u) m_{j;\theta}(x, y, u)}. \quad (9)$$

It is worth noting that (9) can be computed using only first derivatives of α_θ and second derivatives of σ_θ . If N_k is the number of Poisson events between 0 and Δt_k , the parametrix unbiased estimate is then given by

$$\hat{q}_{k+1}(x, y; \zeta_k) = w_{s_{N_k}} m_{k;\theta}(x_{s_{N_k}}, y, t_{k+1} - s_{N_k}),$$

where ζ_k stands for all the randomness required to produce the parametrix estimator (Poisson process and Gaussian random variables).

The stability of this estimator is studied in [Fearnhead et al., 2017] which provides L_p controls for the weight $w_{s_{N_k}}$. The parametrix algorithm mentioned above is a highly flexible procedure to obtain such an unbiased estimate for a much broader class of diffusions than Poisson based estimations which require strong assumptions. However, as the update (9) involves the difference of two Kolmogorov operators, the parametrix estimator of the transition density may be negative, and has no reason to satisfy (8). Thus, the SMC algorithms described above cannot be implemented.

4 Backward importance sampling for PODs

4.1 Positive parametrix estimates

Following [Fearnhead et al., 2010a], Wald's identity for martingales may be used to obtain a new estimator from the parametrix approach, which is guaranteed to be positive. This estimator is defined up to an unknown

constant of proportionality, which is removed when the importance weights are normalized in equation (10). This approach, rather than setting negative weights to 0, which would lead to a biased estimate, uses extra simulation to obtain positiveness. This is done while ensuring that the weights remain unbiased up to a common constant of proportionality. Assume that the distribution \mathbf{K}_k of the additional random variables ζ_k and the estimator \bar{r}_k are obtained with the parametrix estimator.

Particle filtering. For all $k \geq 0$, the Wald-based random weight particle filtering proceeds as follows.

1. For all $1 \leq i \leq N$, sample a new particle as described in Section 3.1.
 - (a) Sample I_k^i in $\{1, \dots, N\}$ with probabilities proportional to $\{\omega_{k-1}^j \vartheta_{k-1}(\xi_{k-1}^j)\}_{1 \leq j \leq N}$.
 - (b) Sample ξ_k^i with distribution $p_{k-1}(\xi_{k-1}^{I_k^i}, \cdot)$.
2. For all $1 \leq i \leq N$, set $\omega_k^i = 0$.
3. While there exists $i_* \in \{1, \dots, N\}$ such that $\omega_k^{i_*} \leq 0$, for all $1 \leq i \leq N$, sample ζ_k^i with distribution $\mathbf{K}_k(\xi_{k-1}^{I_k^i}, \xi_k^i; \cdot)$ (i.e. compute a parametrix estimator of the transition density) and set

$$\omega_k^i = \omega_k^i + \frac{\bar{r}_{k-1}(\xi_{k-1}^{I_k^i}, \xi_k^i; \zeta_k^i)}{\vartheta_{k-1}(\xi_{k-1}^{I_k^i}) p_{k-1}(\xi_{k-1}^{I_k^i}, \xi_k^i)}.$$

Backward simulation. For all $1 \leq i \leq N$, the backward importance sampling step proceeds then as follows.

1. For all $1 \leq j \leq \tilde{N}$, sample $J_{k+1}^{(i,j)}$ in $\{1, \dots, N\}$ with probabilities proportional to $(\omega_k^i)_{i=1}^N$.
2. For all $1 \leq j \leq \tilde{N}$, set $\varpi_k^{(i,j)} = 0$.
3. While there exist $j_* \in \{1, \dots, \tilde{N}\}$ such that $\varpi_k^{(i,j_*)} \leq 0$, for all $1 \leq j \leq \tilde{N}$, sample $\zeta_k^{(i,j)}$ with distribution $\mathbf{K}_k(\xi_k^{J_{k+1}^{(i,j)}}, \xi_{k+1}^i; \cdot)$ and set

$$\varpi_k^{(i,j)} = \varpi_k^{(i,j)} + \bar{r}_k(\xi_k^{J_{k+1}^{(i,j)}}, \xi_{k+1}^i; \zeta_k^{(i,j)}).$$

4.2 AR-free online smoothing

As the positive parametrix-based estimate does not satisfy the upper bound condition of (8), the statistics are updated recursively with an importance sampling step: for all $1 \leq i \leq N$,

$$\tau_{k+1}^i = \sum_{j=1}^{\tilde{N}} \frac{\varpi_k^{(i,j)}}{\mathcal{W}_k^i} \left(\tau_k^{J_{k+1}^{(i,j)}} + \tilde{h}_k \left(\xi_k^{J_{k+1}^{(i,j)}}, \xi_{k+1}^i \right) \right), \quad (10)$$

where $\varpi_k^{(i,j)}$, $1 \leq j \leq \tilde{N}$ are computed using the parametrix estimate combined with Wald's identity. Then, the estimator of the conditional expectation of the additive functional is set as

$$\phi_{0:n|n}^{N,IS}[h_{0:n}] := \sum_{i=1}^N \frac{\omega_n^i}{\Omega_n} \tau_n^i.$$

This estimator does not rely on an accept reject mechanism and is therefore less computationally intensive and can be used under reasonable assumptions for many SDEs. In addition, as shown in Section 6, this does not affect the statistical efficiency of the algorithm.

5 Discussion

This paper proposes a solution to overcome the two main challenges when it comes to perform online smoothing for generic SDEs i.e. obtaining a positive and almost surely bounded estimate of the transition density to run the backward acceptance rejection mechanism.

1. Note that the proposed backward importance sampling may be used to approximate expectations under the smoothing distributions for general state space hidden Markov models and is not restricted to POD processes. As illustrated in Section 6 this approach may lead to significant gains in computational time for similar performance as the acceptance rejection approach.
2. The proposed estimator, unlike the existing methods such as GPE-based algorithms, applies to a large range of multivariate diffusion processes (see [Andersson et al., 2017] and [Fearnhead et al., 2017]).
3. Theoretical guarantees, such as consistency and asymptotic normality of (6), remain to be proved. This should be an extension of [Gloaguen et al., 2019], however this would imply few technicalities which are out of the scope of this paper.
4. The bias of the PaRIS algorithm may be shown to be of order $\mathcal{O}((1 + 1/\tilde{N})/N)$ and vanishes as N goes to infinity for any choice of $\tilde{N} \geq 2$. The exact sampling being replaced by an importance sampling step, the conjecture is that the bias of the proposed algorithm involves a $\mathcal{O}(1/\tilde{N})$ term which does not vanish as N goes to infinity. However, the empirical study illustrates that \tilde{N} may be chosen to increase with N sufficiently slowly to remove the additional bias term while simultaneously ensuring better computational performance. This empirical analysis could be supported by theoretical guarantees and motivates future developments.

6 Numerical experiments

6.1 Sine model

This section investigates the performance of the proposed algorithm to compute expectations under the smoothing distributions in a context where alternatives are available for comparison. Consider the Sine model where $(X_t)_{t \geq 0}$ is assumed to be a weak solution to

$$dX_t = \sin(X_t - \theta)dt + dW_t, \quad X_0 = x_0.$$

This simple model has no explicit transition density, however, a General Poisson estimator which satisfies (8) can be computed by simulating Brownian bridges, (see [Beskos et al., 2006b]). Therefore, the backward importance sampling technique proposed in this paper can be compared to the usual acceptance-rejection algorithm described in Section 3.2. For this simple comparison, observations are received at evenly spaced times $t_0 = 0, \dots, t_{10} = 5$ from the model

$$Y_k = X_{t_k} + \varepsilon_k, \quad 0 \leq k \leq n = 10, \quad (11)$$

where $(\varepsilon_k)_{0 \leq k \leq 10}$ are i.i.d. Gaussian random variables with mean 0 and variance 1. In this experiment $\theta = \pi/4$. The proposal distribution p_k for the particle filtering approximation is chosen as the following approximation of the optimal filter:

$$p_k(x_k, x_{k+1}) \propto q_{k+1}^{\text{Eul}}(x_k, x_{k+1})g_{k+1}(x_{k+1}, Y_{k+1}), \quad (12)$$

where q_{k+1}^{Eul} is the probability density function of Gaussian distribution with mean $\Delta \sin(x_k - \theta)$ and variance Δ where $\Delta = 1/2$, i.e. the Euler approximation of the Sine SDE, and g_k is the probability density function of the law of Y_k given X_{t_k} i.e. of a Gaussian random variable with mean X_{t_k} and variance 1. As the observation model is linear and Gaussian, the proposal distribution is therefore Gaussian with explicit mean and variance.

In this first experiment, particles are used to solve the state estimation problem for the first observation i.e. to compute an estimate of $\mathbb{E}[X_0 | Y_{0:n}]$. Figure 1 displays the computational complexity and the estimation of the posterior mean with the acceptance-rejection algorithm and the proposed backward sampling technique as a function of \tilde{N} . In this setting, $N = 100$, and each unbiased estimate of \hat{q} is computed using 30 Monte Carlo replicates.

For $\tilde{N} = 2$ (which is the recommended value for the PaRIS algorithm, see [Olsson et al., 2017]), our estimate shows a bias, which is no surprise, as it is based on a biased normalized importance sampling step. However, this bias quickly vanishes for $\tilde{N} \geq 10$. Interestingly, our method comes with a drastic (a factor 10) reduction of computational time. The vanishing of the bias might induce more backward sampling, but this remains much faster than the acceptance rejection method with $\tilde{N} = 2$.

Then, the same estimation was performed (on the same data set) for N varying from 50 to 2000. In this context, \tilde{N} was set to 2 for the AR method. To have an empirical intuition of how \tilde{N} must vary with N , the importance sampling algorithm is applied with $\tilde{N} = N^{0.5}$, $N^{0.6}$ and $N/10$ (as this last value was sufficient in the first experiment to avoid any bias).

The results are shown in Figure 2. A small bias might appear for $N = 2000$ and $\tilde{N} = 45$ ($\approx 2000^{0.5}$), but no bias is visible for $N^{0.6}$ and $N/10$. As expected, the gain in time, compared to the state of the art algorithm, remains important (even if it decreases as \tilde{N} increases). It is worth noting that the variance of the computational time is greatly reduced compared to the AR technique.

6.2 Stochastic Lotka-Volterra model

This section sets the focus on a stochastic model describing in continuous time the population dynamics in a predator-prey system, as fully discussed in [Hening and Nguyen, 2018]. The bivariate process $(X_t)_{t \geq 0}$ of predators and preys abundances is assumed to follow the stochastic Lotka-Volterra model, i.e. it is the solution to:

$$dX_t = \alpha_\theta(X_t)dt + \begin{pmatrix} X_1(t) & 0 \\ 0 & X_2(t) \end{pmatrix} \Gamma d\mathbf{W}_t, \quad (13)$$

where \mathbf{W}_t is a vector of independent standard Wiener processes, Γ a 2×2 matrix, and for $x = (x_1, x_2)^T$:

$$\alpha_\theta(x) = \begin{pmatrix} x_1(a_{10} - a_{11}x_1 - a_{12}x_2) \\ x_2(-a_{20} + a_{21}x_1 - a_{22}x_2) \end{pmatrix}.$$

In this context, the unknown parameter to be estimated is

$$\theta = (a_{10}, a_{11}, a_{12}, a_{20}, a_{21}, a_{22}, \Gamma).$$

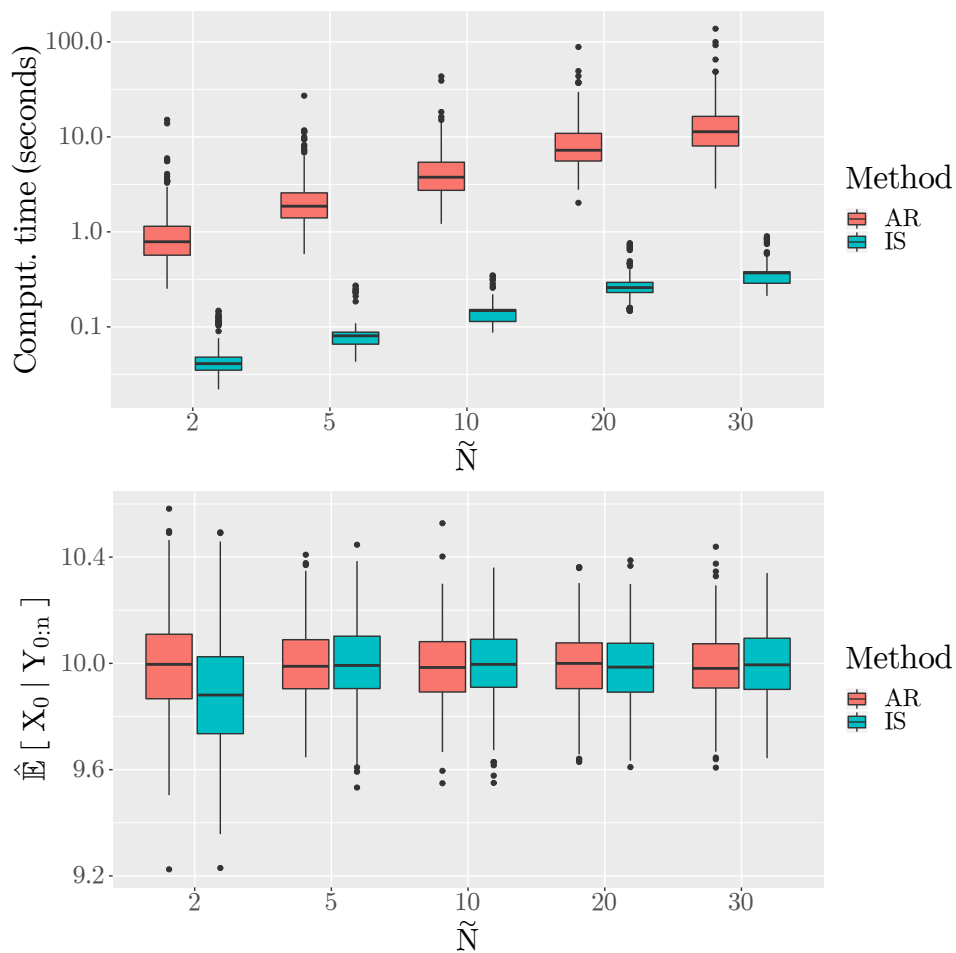


Figure 1: Computational complexity and estimation of a posterior mean as a function of the number of backward samples. Results are shown for the state of the art acceptance-rejection algorithm and the proposed backward importance sampling technique.

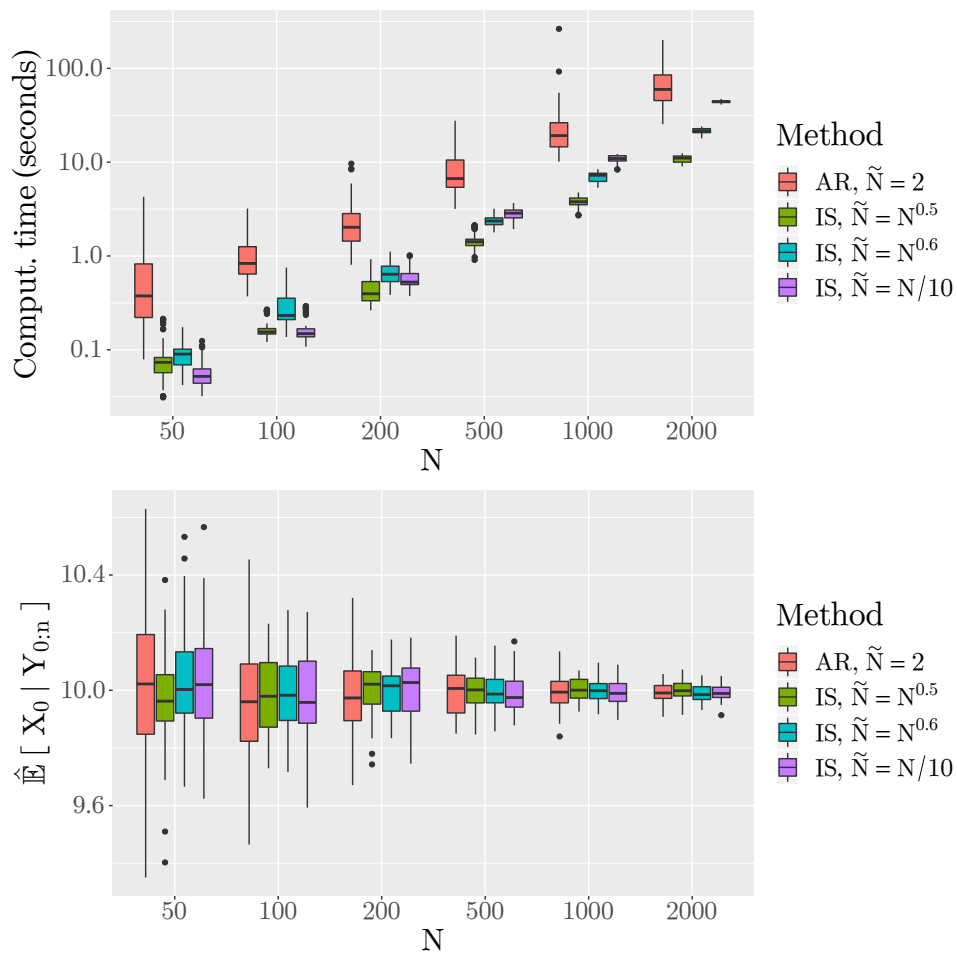


Figure 2: Computational complexity and estimation of a posterior mean as a function of the number of particles. Results are shown for the state of the art acceptance-rejection algorithm and the proposed backward importance sampling technique. The number of backward samples is set to 2 for the AR, and $N/10$ for the IS.

The observation model follows a widespread framework in ecology where the abundance of preys and predators are observed through some abundance index at discrete times t_0, \dots, t_n such that:

$$Y_{t_k} = \begin{pmatrix} c_1 X_1(t_k) e^{\epsilon_{t_k}^{(1)}} \\ c_2 X_2(t_k) e^{\epsilon_{t_k}^{(2)}} \end{pmatrix}, \quad (14)$$

where $c = (c_1, c_2)^T$ is a known parameter (the observed fraction of the population) and $\{\epsilon_{t_k} = (\epsilon_{t_k}^{(1)}, \epsilon_{t_k}^{(2)})\}_{1 \leq k \leq n}$ are i.i.d. random variables distributed as a $\mathcal{N}_2(-\frac{1}{2} \text{diag } \Sigma, \Sigma)$ where Σ is an unknown 2×2 covariance matrix.

It is straightforward to show that for a generic θ , in the SDE defined by (13), the drift function cannot be written (even after the Lamperti transform) as the gradient of a potential. Therefore, the General Poisson estimator cannot be used as an unbiased estimator of the transition density, and the method proposed in this paper is the only solution to obtain a consistent estimate of the target expectations.

The proposal distribution for the particle filter is again a trade off between model dynamics and the observation model (full details are given in the appendix). The simulated set of particles is used to obtain estimates of the true abundances given the observations, both on synthetic and real data.

Synthetic data

In a first approach, simulated data are obtained from the model given by (13) and (14) for a known set of parameters. Chosen values of θ , Σ , c_1 and c_2 for the experiment are given in the appendix. The model is used to simulate abundances indexes Y_0, \dots, Y_{300} at times $t_0 = 0, \dots, t_{300} = 3$. The associated time series (after a division by c) is shown in Figure 3 (left panel).

In this experiment, the goal is to obtain an estimate of the actual predator-prey abundances given all the observed abundances indexes $Y_{0:n}$. Our estimate is given by the set of conditional expectations $\{\mathbb{E}[X_k | Y_{0:n}]\}_{k=0, \dots, n}$, approximated using our backward importance sampling PaRIS smoother, which is run using the true parameters. Figure 3 shows the estimated abundance trajectory over time. The proposed algorithm manages to estimate efficiently the actual abundance from noisy data and a model with an intractable transition density.

Hares and lynx data

In this section, the model defined by equations (13) and (14) is applied to the Hudson Bay company data, giving the number of hares and lynx trapped in Canada during the first 20 years of the 20th century (available in [Odum and Barrett, 1971]).

As parameters are unknown in this case, maximum likelihood inference is performed using an EM [Dempster et al., 1977] algorithm to obtain an estimate $\hat{\theta}$. As explained in the introduction, it is then required to estimate iteratively, from an initial guess θ_0 , the conditional expectation given in equation (2). This E step is performed using the particle smoother introduced in this paper. At each iteration, the estimator θ_k is updated by finding a parameter θ_{k+1} for which $\hat{Q}(\theta_{k+1}, \theta_k) > \hat{Q}(\theta_k, \theta_k)$, with a gradient free evolution strategy [Hansen, 2006].

The last estimate $\hat{\theta}$ obtained with this EM algorithm is used to estimate the actual abundances in the model (similarly to the synthetic data case). Figure 4 shows estimates of $\mathbb{E}_{\hat{\theta}}[X_k | Y_{0:n}]$ obtained with 30 independent runs of our algorithm. The particle smoother is implemented using $N = 200$ particles and $\tilde{N} = 20$. The replicates show that the variance of our estimator (for a given set of observations) is much smaller than the one of the *poor man smoother*. This algorithm approximates the smoothing distributions at

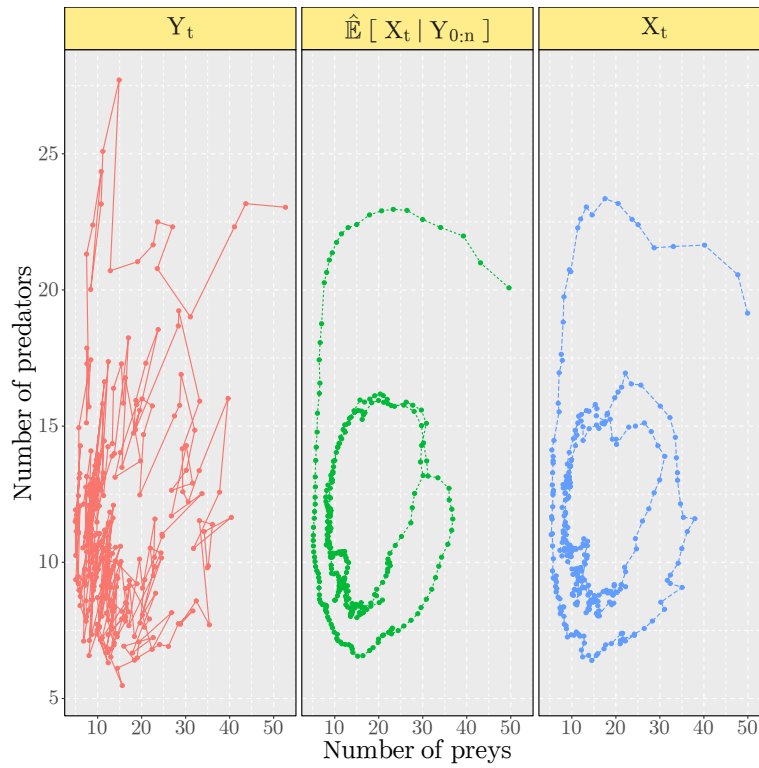


Figure 3: Estimated predator-prey abundances (center) in a stochastic Lotka Volterra model using our backward sampling estimate on simulated abundance indexes (left). Right panel shows the ground truth.




Figure 4: Estimated Hares-Lynx abundances using the Hudson bay company data set together with a stochastic Lotka Volterra model. Both our IS-PaRIS smoother and the poor man smoother are performed to approximate the MLE and solve the tracking problem. Blue crosses show the observations.

time n by the weighted samples $\{(\xi_{0:n}^\ell, \omega_n^\ell)\}_{\ell=1\dots N}$ where the particle trajectories $\xi_{0:n}^\ell$ are obtained using the ancestral line of each last sample: for all $1 \leq k \leq n$,

$$\xi_{0:k}^\ell = (\xi_{0:k-1}^\ell, \xi_k^\ell).$$

The variance of the estimates based on this ancestor tree is doomed to failure due to the degeneracy caused by the successive resampling steps.

References

- [Aït-Shalia, 2008] Aït-Shalia, Y. (2008). Closed-form likelihood expansions for multivariate diffusions. *Annals of Statistics*, 36(2):906–937.
- [Andersson et al., 2017] Andersson, P., Kohatsu-Higa, A., et al. (2017). Unbiased simulation of stochastic differential equations using parametrix expansions. *Bernoulli*, 23(3):2028–2057.

- [Beskos et al., 2006a] Beskos, A., Papaspiliopoulos, O., and Roberts, G. O. (2006a). Retrospective exact simulation of diffusion sample paths with applications. *Bernoulli*, 12(6):1077–1098.
- [Beskos et al., 2006b] Beskos, A., Papaspiliopoulos, O., Roberts, G. O., and Fearnhead, P. (2006b). Exact and computationally efficient likelihood-based estimation for discretely observed diffusion processes (with discussion). *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(3):333–382.
- [Briers et al., 2010] Briers, M., Doucet, A., and Maskell, S. (2010). Smoothing algorithms for state–space models. *Annals of the Institute of Statistical Mathematics*, 62(1):61.
- [Candanedo et al., 2017] Candanedo, L. M., Feldheim, V., and Deramaix, D. (2017). A methodology based on hidden markov models for occupancy detection and a case study in a low energy residential building. *Energy and Buildings*, 148:327–341.
- [Cappé et al., 2005] Cappé, O., Moulines, E., and Rydén, T. (2005). *Inference in Hidden Markov Models*. Springer.
- [Del Moral et al., 2010] Del Moral, P., Doucet, A., and Singh, S. S. (2010). A backward particle interpretation of feynman-kac formulae. *ESAIM: Mathematical Modelling and Numerical Analysis*, 44(5):947–975.
- [Dempster et al., 1977] Dempster, A., Laird, N., and Rubin, D. (1977). Maximum likelihood from incomplete data via the EM algorithm (with discussion). 39:1–38.
- [Douc et al., 2011] Douc, R., Garivier, A., Moulines, E., and Olsson, J. (2011). Sequential monte carlo smoothing for general state space hidden markov models. *The Annals of Applied Probability*, 21(6):2109–2145.
- [Douc et al., 2014] Douc, R., Moulines, E., and Stoffer, D. (2014). *Nonlinear time series: theory, methods and applications with R examples*. CRC Press.
- [Douc et al., 2013] Doucet, A., de Freitas, N., and Gordon, N. (2013). *Sequential Monte Carlo methods in practice*. Springer Science & Business Media.
- [Douc et al., 2000] Doucet, A., Godsill, S., and Andrieu, C. (2000). On sequential monte carlo sampling methods for bayesian filtering. *Statistics and computing*, 10(3):197–208.
- [Dubarry and Le Corff, 2011] Dubarry, C. and Le Corff, S. (2011). Fast computation of smoothed additive functionals in general state-space models. *2011 IEEE Statistical Signal Processing Workshop (SSP)*, pages 197–200.
- [Dubarry and Le Corff, 2013] Dubarry, C. and Le Corff, S. (2013). Non-asymptotic deviation inequalities for smoothed additive functionals in nonlinear state-space models. *Bernoulli*, 19(5B):2222–2249.
- [Fearnhead et al., 2017] Fearnhead, P., Latuszynski, K., Roberts, G. O., and Sermaidis, G. (2017). Continuous-time importance sampling: Monte carlo methods which avoid time-discretisation error. *arXiv preprint arXiv:1712.06201*.
- [Fearnhead et al., 2010a] Fearnhead, P., Papaspiliopoulos, O., Roberts, G., and Stuart, A. (2010a). Random weight particle filtering of continuous time stochastic processes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 72(4):497–512.

- [Fearnhead et al., 2008] Fearnhead, P., Papaspiliopoulos, O., and Roberts, G. O. (2008). Particle filters for partially observed diffusions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 70(4):755–777.
- [Fearnhead et al., 2010b] Fearnhead, P., Wyncoll, D., and Tawn, J. (2010b). A sequential smoothing algorithm with linear computational cost. *Biometrika*, 97(2):447–464.
- [Gassiat et al., 2016] Gassiat, É., Cleynen, A., and Robin, S. (2016). Inference in finite state space non parametric hidden markov models and applications. *Statistics and Computing*, 26(1-2):61–71.
- [Gerber and Chopin, 2017] Gerber, M. and Chopin, N. (2017). Convergence of sequential quasi-Monte Carlo smoothing algorithms. *Bernoulli*, 23(4B):2951–2987.
- [Gloaguen et al., 2018] Gloaguen, P., Etienne, M.-P., and Le Corff, S. (2018). Online sequential monte carlo smoother for partially observed diffusion processes. *EURASIP Journal on Advances in Signal Processing*, 2018(1):9.
- [Gloaguen et al., 2019] Gloaguen, P., Le Corff, S., and Olsson, J. (2019). Pseudo marginal sequential Monte Carlo methods for general state spaces. Applications to recursive maximum likelihood. *arXiv:1908.07254*.
- [Gordon et al., 1993] Gordon, N. J., Salmond, D. J., and Smith, A. F. (1993). Novel approach to nonlinear/non-gaussian bayesian state estimation. In *IEE proceedings F (radar and signal processing)*, volume 140, pages 107–113. IET.
- [Hansen, 2006] Hansen, N. (2006). The cma evolution strategy: a comparing review. In *Towards a new evolutionary computation*, pages 75–102. Springer.
- [Hening and Nguyen, 2018] Hening, A. and Nguyen, D. H. (2018). Persistence in stochastic lotka–volterra food chains with intraspecific competition. *Bulletin of mathematical biology*, 80(10):2527–2560.
- [Michelot et al., 2016] Michelot, T., Langrock, R., and Patterson, T. A. (2016). movehmm: an r package for the statistical modelling of animal movement data using hidden markov models. *Methods in Ecology and Evolution*, 7(11):1308–1315.
- [Nguyen et al., 2017] Nguyen, T., Le Corff, S., and Moulines, É. (2017). On the two-filter approximations of marginal smoothing distributions in general state-space models. *Advances in Applied Probability*, 50(1):154–177.
- [Odum and Barrett, 1971] Odum, E. P. and Barrett, G. W. (1971). *Fundamentals of ecology*, volume 3. Saunders Philadelphia.
- [Olsson et al., 2011] Olsson, J., Ströjby, J., et al. (2011). Particle-based likelihood inference in partially observed diffusion processes using generalised poisson estimators. *Electronic Journal of Statistics*, 5:1090–1122.
- [Olsson et al., 2017] Olsson, J., Westerborn, J., et al. (2017). Efficient particle-based online smoothing in general hidden markov models: the paris algorithm. *Bernoulli*, 23(3):1951–1996.
- [Pitt and Shephard, 1999] Pitt, M. K. and Shephard, N. (1999). Filtering via simulation: Auxiliary particle filters. *Journal of the American statistical association*, 94(446):590–599.

- [Rabiner, 1989] Rabiner, L. (1989). A tutorial on hidden markov models and selected applications in speech recognition. In *Proceedings of the IEEE*, pages 257–286.
- [Särkkä, 2013] Särkkä, S. (2013). *Bayesian Filtering and Smoothing*. Cambridge University Press, New York, NY, USA.
- [Särkkä et al., 2007] Särkkä, S., Vehtari, A., and Lampinen, J. (2007). Rao-Blackwellized particle filter for multiple target tracking. *Information Fusion*, 8(1):2–15.
- [Wang et al., 2017] Wang, X., Lebarbier, E., Aubert, J., and Robin, S. (2017). Variational inference for coupled hidden markov models applied to the joint detection of copy number variations. *The International Journal of Biostatistics*, 15.
- [Yau et al., 2011] Yau, C., Papaspiliopoulos, O., Roberts, G. O., and Holmes, C. (2011). Bayesian non-parametric hidden Markov models with applications in genomics. 73:1–21.
- [Zucchini et al., 2017] Zucchini, W., Mac Donald, I., and Langrock, R. (2017). *Hidden Markov models for time series: an introduction using R*. CRC Press.