



# Particle-based online estimation of tangent filters with application to parameter estimation in nonlinear state-space models

Jimmy Olsson<sup>1</sup> · Johan Westerborn Alenlöv<sup>1</sup>

Received: 22 December 2017 / Revised: 18 October 2018 / Published online: 19 November 2018  
© The Institute of Statistical Mathematics, Tokyo 2018

## Abstract

This paper presents a novel algorithm for efficient online estimation of the filter derivatives in general hidden Markov models. The algorithm, which has a linear computational complexity and very limited memory requirements, is furnished with a number of convergence results, including a central limit theorem with an asymptotic variance that can be shown to be uniformly bounded in time. Using the proposed filter derivative estimator, we design a recursive maximum likelihood algorithm updating the parameters according to the gradient of the one-step predictor log-likelihood. The efficiency of this online parameter estimation scheme is illustrated in a simulation study.

**Keywords** Parameter estimation · Recursive maximum likelihood · State-space models · Tangent filter · Sequential Monte Carlo methods · Central limit theorem · Particle filters

## 1 Introduction

The general state-space *hidden Markov models* form a powerful statistical modelling tool which is presently applied across a wide range of scientific and engineering disciplines; see e.g. Cappé (2001), Cappé et al. (2005) and the references therein for such examples. In the literature, the term “hidden Markov model” is often restricted to the case where the state space of the hidden chain is a finite set, and we hence use the term “state-space model” (SSM) to stress that the state spaces of the models that we consider are completely general. More specifically, SSMs are bivariate stochastic

---

✉ Johan Westerborn Alenlöv  
johawes@math.kth.se

Jimmy Olsson  
jimmyol@math.kth.se

<sup>1</sup> Department of Mathematics, KTH Royal Institute of Technology, 100 44 Stockholm, Sweden

processes consisting of an observable process  $\{Y_t\}_{t \in \mathbb{N}}$  and an unobservable Markov chain  $\{X_t\}_{t \in \mathbb{N}}$ , referred to as the *state process* and *observation process* and taking on values in some general state spaces  $Y$  and  $X$ , respectively. When using these models in practice, the statistician is typically interested in calculating conditional distributions of the unobservable states given some fixed observation record or data-stream  $\{y_t\}_{t \in \mathbb{N}}$ . Generally, these distributions cannot be expressed in a closed form, and the user needs to rely on approximations. Moreover, any SSM is typically parameterised by unknown model parameters  $\theta$  which need to be calibrated *ex ante*.

The frequentist approach to parameter estimation consists in maximising the *likelihood function*  $\theta \mapsto L_\theta(y_{0:t})$ , where  $L_\theta(y_{0:t})$  denotes the joint density of the observations evaluated at the given data  $y_{0:t} = (y_0, \dots, y_t)$  (this will be our generic notation for vectors), or, equivalently, the *log-likelihood function*  $\theta \mapsto \ell_\theta(y_{0:t}) := \log L_\theta(y_{0:t})$ . Even though the likelihood cannot generally be expressed in a closed form, there are several approaches to approximate maximisation of the same. These methods are typically based on either the *expectation–maximisation* algorithm (Dempster et al. 1977), where the maximisation is performed over an intermediate quantity, or the tools of *gradient-based optimisation*. The latter algorithms produce, in a *steepest ascent* manner, a sequence  $\{\theta_n\}_{n \in \mathbb{N}}$  of parameter estimates converging the maximum likelihood estimator; i.e. at each iteration, a parameter estimate  $\theta_n$  is updated recursively through a move in the direction given by an estimate of the *score function*  $\nabla_\theta \ell_{\theta_n}(y_{0:t})$ .

There are mainly two approaches to estimation of the score function. The first one goes via *Fisher's identity* (see e.g. Cappé et al. 2005, Proposition 10.1.6), which relates the score function of the observed data to that of the complete data and thus transfers the problem to that of *joint smoothing*, i.e. the computation of the joint posterior distribution of  $X_{0:t}$  given  $Y_{0:t} = y_{0:t}$ . The second approach, which is the focus of the present paper, goes via the decomposition

$$\ell_\theta(y_{0:t}) = \sum_{s=0}^t \ell_\theta(y_s \mid y_{0:s-1}), \quad (1)$$

where each *one-step predictor likelihood*  $\ell_\theta(y_s \mid y_{0:s-1})$  (with  $\ell_\theta(y_0 \mid y_{0:-1}) := \ell_\theta(y_0)$ ) can be computed via the *sensitivity equations*; see Cappé et al. (2005, Section 10.2.4) and Sect. 4 for details. Let  $\pi_{s;\theta}(x_s)$  be the density of the *predictor* at time  $s$ , i.e. the distribution of the state  $X_s$  conditional to the observation record  $Y_{0:s-1} = y_{0:s-1}$ ; it is then, by swapping the order of the nabla and integration operations, possible to express the gradient  $\nabla_\theta \ell_\theta(y_s \mid y_{0:s-1})$  as a functional of the gradient  $\nabla_\theta \pi_{s;\theta}(x_s)$ , the latter being typically referred to as the *filter derivative* or *tangent filter* or *filter sensitivity*. It should be noted that despite its name, the tangent filter is the gradient of the *predictor* rather than the *filter*, where the filter at time  $s$  is defined as the distribution of  $X_s$  conditionally to  $Y_{0:s} = y_{0:s}$ , i.e. the updated predictor.

Appealingly, the recursive updating scheme for the tangent filters given by the sensitivity equations requires only access to the filter distributions (and not the full joint-smoothing distribution, as in the case of Fisher identity-based score approximation). If estimation is to be carried through in the *offline* mode, by processing repeatedly

a fixed batch of observations until convergence, the approach based on Fisher's identity is typically computationally more efficient and hence to be preferred. On the other hand, if estimation is to be carried through in an *online* manner, i.e. through a single sweep of a (potentially infinitely) long sequence of observations, then the sensitivity equation approach is better suited, as the filter distributions can be updated recursively at constant computational cost. Online parameter learning is relevant when dealing with very big batches of data or in scenarios requiring the parameters to be continuously updated on the basis of a continuous stream of arriving observations.

If the state space of the latent Markov chain is a finite set or if the model belongs to the linear Gaussian SSMs, exact computation of the filter distributions—and thus the filter derivatives and the gradient of the one-step predictor log-likelihood—is possible (Le Gland and Mevel 1997; Cappé et al. 2005). However, as touched upon above, the general case calls for approximation of these quantities, and existing approaches to nonlinear filtering can typically be divided into two classes. The first class contains methods relying on linearisation, such as the *extended Kalman filter*, the *unscented Kalman filter* (Anderson and Moore 1979; Julier and Uhlmann 1997), etc., whose success is generally limited in the presence of significant nonlinear or non-Gaussian model components. The second class is formed by the *sequential Monte Carlo* (SMC) or *particle filtering* approaches (Gordon et al. 1993), where a sample of particles with associated importance weights is propagated sequentially and randomly in time according to the model dynamics in order to approximate the flows of filter and predictor distributions. Due to their high capability of solving potentially very difficult nonlinear/non-Gaussian filtering problems, SMC methods have, as indicated by the over 8000 Google Scholar citations of the first monograph Doucet et al. (2001) on the topic, seen a rapid and ever-increasing interest during the last decades.

## 1.1 Previous work

For an exposé of current particle-based approaches to parameter estimation in nonlinear SSMs, we refer to Kantas et al. (2015). SMC-based score function estimation via tangent filters was discussed by Doucet and Tadić (2003), Poyiadjis et al. (2005, 2011), Del Moral et al. (2015).

The work Del Moral et al. (2015), which can be viewed as the state of the art when it concerns SMC-based tangent filter estimation, expresses each tangent filter as a centred, smoothed expectation of an additive state functional. Using the so-called *backward decomposition* of the joint-smoothing distribution and a recursive form of this decomposition going back to Cappé (2011), Del Moral et al. (2010) are able to form, on-the-fly and through the perspicacious propagation of a number of smoothed statistics, one statistic per particle, non-collapsed particle approximations of such additive smoothed expectations. Moreover, Del Moral et al. (2010) furnish their estimator with a central limit theorem (CLT) describing the weak convergence, as the size of the particle sample tends to infinity, of the particle tangent filter approximations to the exact counterparts. Appealingly, the fact that the smoothed expectations are centred allows, under strong mixing assumptions, the asymptotic variance of the mentioned CLT to be bounded uniformly in time, which establishes the long-term stochastic sta-

bility of the algorithm for large sample sizes. However, since the propagation of the smoothed statistic associated with a given particle requires an expectation under the ancestral posterior of that particle to be computed by summation, the computational complexity of the algorithm is *quadratic* in the number of particles. The user is hence forced to keep the number of particles at a modest value, which may imply severe numerical instability in practice (see Sect. 4.3 for an illustration).

## 1.2 Our contribution

On the basis of the approach taken by [Del Moral et al. \(2015\)](#), we develop an efficient and numerically stable algorithm for particle approximation of the filter derivatives in SSMs. The approach that we propose includes, as a sub-routine, the *particle-based, rapid incremental smoother* (PaRIS) (introduced recently by us in [Olsson and West-erborn 2017](#)), an algorithm that is tailor-made for online approximation of smoothed additive state functionals.

By replacing, in the forward recursion proposed by [Del Moral et al. \(2015\)](#), the computation of each smoothed statistic by a conditionally unbiased Monte Carlo estimate, the computational complexity becomes, under mild assumptions, *linear* in the number of particles. Since the original estimator can be viewed as a Rao–Blackwellisation of the new one, we will sometimes refer to this measure as “*de-Rao–Blackwellisation*”. This allows the user to mount a considerably larger amount of particles for a given computational budget, which compensates by far for the slight increase in asymptotic variance added by the additional sampling step. Thus, the technique that we propose has some similarities with that used in particle-based *two-filter smoothing*, where a linear computational complexity is obtained by replacing certain marginalisations by random sampling (see [Fearnhead et al. 2010](#)) (however, for the sake of completeness we remark that two-filter particle smoothers, which require a so-called *backward information filter* to be run backwards in time, cannot be used in online settings). By reusing techniques developed by [Olsson and Westerborn \(2017\)](#) and [Del Moral et al. \(2015\)](#), we are able to furnish the proposed algorithm with an exponential concentration inequality and a CLT, whose asymptotic variance is given by that of the Rao–Blackwellised counterpart plus one additional term being inversely proportional to sample size of the supplementary, PaRISian sampling operation. This result is analogous to the results obtained by [Nguyen et al. \(2017\)](#) for the case of particle-based two-filter smoothing. Moreover, by assuming strong mixing of the model, the asymptotic variance can be bounded uniformly in time.

In the second part of the paper, we cast our PaRIS-based tangent filter estimator into the framework of recursive maximum likelihood estimation. The numerical performance of this online parameter estimation algorithm is investigated in a simulation study, where we are able to report significant improvement over the approach of [Del Moral et al. \(2015\)](#) for online parameter learning in the stochastic volatility model of [Hull and White \(1987\)](#). In addition, we propose our algorithm as a prototypical solution to the *simultaneous localisation and mapping* (SLAM) problem, with promising results.

Finally, we remark that the parameter estimation algorithm was outlined by us in the conference note (Olsson and Westerborn 2016). The present paper provides rigorous theoretical results underpinning the initial observations made by Olsson and Westerborn (2016) and complement the simulation study of that work with additional examples.

### 1.3 Outline

After having introduced, prefatorily, general SSMs, tangent filters, and SMC-based smoothing (including the PaRIS) are presented in Sect. 2, and Sect. 3 presents the proposed approach to tangent filter estimation. In addition, Sect. 3.1 includes theoretical convergence results in the form of an exponential concentration inequality and a CLT with a time-uniform bound on the asymptotic variance. Section 4 applies the PaRISian tangent filter estimator to recursive maximum likelihood estimation, and the resulting online parameter estimation algorithm is tested numerically in Sect. 4.3. Finally, some conclusions are drawn in Sect. 5 and “Appendix A” contains all proofs.

## 2 Preliminaries

### 2.1 General state-space models

In the following, we let  $\mathbb{N}$  denote the natural numbers and set  $\mathbb{N}^* := \mathbb{N} \setminus \{0\}$ . For any measurable space  $(E, \mathcal{E})$ , where  $\mathcal{E}$  is a countably generated  $\sigma$ -algebra, we denote by  $F(\mathcal{E})$  the set of bounded  $\mathcal{E}/\mathcal{B}(\mathbb{R})$ -measurable functions on  $E$ , where  $\mathcal{B}(\mathbb{R})$  denotes the Borel  $\sigma$ -algebra. For any function  $h \in F(\mathcal{E})$ , we let  $\|h\|_\infty := \sup_{x \in E} |h(x)|$  denote the supremum norm of  $h$ . Let  $M(\mathcal{E})$  denote the set of  $\sigma$ -finite measures on  $\mathcal{E}$  and  $M_1(\mathcal{E}) \subset M(\mathcal{E})$  the set of probability measures on the same space.

Now, let  $(X, \mathcal{X})$  and  $(Y, \mathcal{Y})$  be measurable spaces and  $\mathbf{Q}_\theta : X \times \mathcal{X} \rightarrow [0, 1]$  and  $\mathbf{G}_\theta : X \times \mathcal{Y} \rightarrow [0, 1]$  Markov transition kernels. We will often subject kernels to operations such as multiplication and tensor multiplication. Moreover, a transition kernel induces two different operators, one on bounded measurable functions and one on measures; we refer to “Appendix B” for details. Moreover, let  $\chi \in M_1(\mathcal{X})$ . The kernels above are parameterised by a model parameter vector  $\theta$  belonging to some space  $\Theta \subseteq \mathbb{R}^d$ ,  $d \in \mathbb{N}^*$ , and we will use subindices to stress the dependence of any distributions on these parameters. We define an SSM as the canonical Markov chain  $\{(X_t, Y_t)\}_{t \in \mathbb{N}}$  with initial distribution  $\chi \otimes \mathbf{G}_\theta$  (see “Appendix B” for a definition of the product  $\otimes$  of a measure and a kernel) and Markov transition kernel

$$(X \times Y) \times (\mathcal{X} \otimes \mathcal{Y}) \ni ((x, y), A) \mapsto \int \mathbf{1}_A(x', y') \mathbf{Q}_\theta(x, dx') \mathbf{G}_\theta(x', dy'). \quad (2)$$

In this setting, the state process  $\{X_t\}_{t \in \mathbb{N}}$  is assumed to be only partially observed through the observation process  $\{Y_t\}_{t \in \mathbb{N}}$ . Under the dynamics (2),

- (i) the state process  $\{X_t\}_{t \in \mathbb{N}}$  is itself a Markov chain with transition kernel  $\mathbf{Q}_\theta$  and initial distribution  $\chi$ .
- (ii) the observations are, conditionally on the states, independent and such that the conditional distribution of each  $Y_t$  depends on the corresponding  $X_t$  only and is given by the *emission distribution*  $\mathbf{G}_\theta(X_t, \cdot)$ .

For simplicity, we have assumed that the initial distribution  $\chi$  does not depend on  $\theta$ . Throughout this paper, we will consider the fully dominated case where for all  $\theta \in \Theta$ ,  $\mathbf{Q}_\theta$  and  $\mathbf{G}_\theta$  admit densities  $q_\theta$  and  $g_\theta$  with respect to some reference measures  $\mu \in \mathcal{M}(\mathcal{X})$  and  $\nu \in \mathcal{M}(\mathcal{Y})$ , respectively. This means that for all  $\theta \in \Theta$ ,  $x \in \mathcal{X}$ ,  $f \in \mathcal{F}(\mathcal{X})$ , and  $h \in \mathcal{F}(\mathcal{Y})$ ,

$$\begin{aligned} \mathbf{Q}_\theta f(x) &= \int f(x')q_\theta(x, x') \mu(dx'), \\ \mathbf{G}_\theta h(x) &= \int h(y)g_\theta(x, y) \nu(dy). \end{aligned}$$

Likewise,  $\chi$  is assumed to have a density, which we denote by the same symbol, with respect to the same reference measure  $\mu$ .

In the following, we assume that we are given a fixed sequence  $\{y_t\}_{t \in \mathbb{N}}$  of observations of  $\{Y_t\}_{t \in \mathbb{N}}$ , and in order to avoid notational overload, we will generally suppress any dependence on these observations from the notation by denoting, e.g. by  $g_{t;\theta}$  the function  $X \ni x \mapsto g_\theta(x, y_t)$ . For any triple  $(s, s', t) \in \mathbb{N}^3$  such that  $s \leq s'$  we denote by  $\phi_{s:s'|t;\theta}$  the conditional distribution of  $X_{s:s'}$  conditioned on  $Y_{0:t}$ . For any  $f \in \mathcal{F}(\mathcal{X}^{\otimes(s'-s+1)})$ , this distribution can be expressed as:

$$\begin{aligned} \phi_{s:s'|t;\theta} f &= L_\theta(y_{0:t})^{-1} \int \cdots \int f(x_{s:s'}) \left( \prod_{m=0}^t g_{m;\theta}(x_m) \right) \\ &\quad \times \left( \chi(dx_0) \prod_{\ell=0}^{(t \vee s')-1} \mathbf{Q}_\theta(x_\ell, dx_{\ell+1}) \right), \end{aligned}$$

where

$$L_\theta(y_{0:t}) = \int \cdots \int g_{0;\theta}(x_0) \chi(dx_0) \prod_{\ell=0}^{t-1} g_{\ell+1;\theta}(x_{\ell+1}) \mathbf{Q}_\theta(x_\ell, dx_{\ell+1}) \tag{3}$$

is the observed data likelihood. In the cases  $s = s' = t$  and  $s = s' = t + 1$ , we let  $\phi_{t;\theta}$  and  $\pi_{t+1;\theta}$  be shorthand notation for the filter and predictor distributions  $\phi_{t:t|t;\theta}$  and  $\phi_{t+1:t+1|t;\theta}$ , respectively. In the case  $s = 0$  and  $s' = t$ , the distribution  $\phi_{0:t|t;\theta}$  is referred to as the joint-smoothing distribution. The filter and predictor distributions are closely related; indeed, the well-known *filtering recursion* provides that for all  $t \in \mathbb{N}$  and  $f \in \mathcal{F}(\mathcal{X})$ ,

$$\phi_{t;\theta} f = \frac{\pi_{t;\theta}(g_{t;\theta} f)}{\pi_{t;\theta} g_{t;\theta}}, \tag{4}$$

$$\pi_{t+1;\theta} f = \phi_{t;\theta} \mathbf{Q}_\theta f, \tag{5}$$

where, by convention,  $\pi_{0;\theta} := \chi$ .

We will often deal with sums and products on functions with possibly different arguments. Since these functions will be defined on products of  $X$ , we will, when needed, let, with slight abuse of notation, subscripts define the domains of such sums and products. For instance,  $f_t \tilde{f}_t : X \ni x_t \mapsto f_t(x_t) \tilde{f}_t(x_t)$ , while  $f_t + \tilde{f}_{t+1} : X^2 \ni x_{t:t+1} \mapsto f_t(x_t) + \tilde{f}_{t+1}(x_{t+1})$ .

It may be shown (see e.g. Cappé et al. 2005, Proposition 3.3.6) that the state process has still the Markov property when evolving conditionally on  $Y_{0:t} = y_{0:t}$  in the time-reversed direction. Moreover, the distribution of  $X_s$  given  $X_{s+1}$  and  $Y_{0:t} = y_{0:t}$  is, for  $s \leq t$ , given by the backward kernel denoted by  $\overleftarrow{\mathbf{Q}}_{\phi_{s;\theta}}$ . Under the assumption of a fully dominated model, the backward kernel can, for  $x_{s+1} \in X$  and  $f \in F(\mathcal{X})$ , be expressed as:

$$\overleftarrow{\mathbf{Q}}_{\phi_{s;\theta}} f(x_{s+1}) := \frac{\int f(x_s) q_\theta(x_s, x_{s+1}) \phi_{s;\theta}(dx_s)}{\int q_\theta(x'_s, x_{s+1}) \phi_{s;\theta}(dx'_s)}.$$

Consequently, using the backward kernel, we may express the joint-smoothing distribution  $\phi_{0:t|t;\theta}$  as

$$\phi_{0:t|t;\theta} = \phi_{t;\theta} \mathbf{T}_{t;\theta}, \tag{6}$$

where we have defined the kernels

$$\mathbf{T}_{t;\theta} := \begin{cases} \overleftarrow{\mathbf{Q}}_{\phi_{t-1;\theta}} \otimes \overleftarrow{\mathbf{Q}}_{\phi_{t-2;\theta}} \otimes \dots \otimes \overleftarrow{\mathbf{Q}}_{\phi_{0;\theta}} & \text{for } t \in \mathbb{N}^*, \\ \text{id} & \text{for } t = 0, \end{cases}$$

where  $\otimes$  denotes kernel tensor product; see ‘‘Appendix B’’ for details. Note that for all  $x \in X$ , the distribution  $\mathbf{T}_{t;\theta}(x, \cdot)$  on the product space  $\mathcal{X}^{\otimes t}$  describes the law of the inhomogeneous Markov chain initialised at  $x \in X$  and evolving backwards in time according to the backward kernels.

As we will see in the following, the modeller is often required to compute smoothed expectations of additive objective functions of type

$$h_t(x_{0:t}) := \sum_{s=0}^{t-1} \tilde{h}_s(x_{s:s+1}),$$

where for all  $0 \leq s \leq t - 1$ ,  $\tilde{h}_s$  is a measurable function on  $X^2$ . This setting allows for recursive computation of  $\{\mathbf{T}_{t;\theta} h_t\}_{t \in \mathbb{N}}$  according to

$$\mathbf{T}_{t+1;\theta} h_{t+1}(x_{t+1}) = \int \{\mathbf{T}_{t;\theta} h_t(x_t) + \tilde{h}_t(x_{t:t+1})\} \overleftarrow{\mathbf{Q}}_{\phi_{t;\theta}}(x_{t+1}, dx_t); \tag{7}$$

see Cappé (2011), Del Moral et al. (2010). The recursion (7) will be a key ingredient also in the coming developments.

### 2.2 Tangent filters

The gradient of the prediction filter, known as the filter derivative or the tangent filter, is a signed measure which we, following (Del Moral et al. 2015, Section 2), define as follows. We start with considering the gradient of  $\phi_{0:t|t-1;\theta}$  and then extract the marginal of this quantity. First, note that for all  $f_{0:t} \in F(\mathcal{X}^{\otimes(t+1)})$ ,

$$\begin{aligned} & \frac{1}{L_\theta(y_{0:t-1})} \nabla_\theta \int f_{0:t}(x_{0:t}) \chi(dx_0) \prod_{s=0}^{t-1} g_{s;\theta}(x_s) q_\theta(x_s, x_{s+1}) \mu^{\otimes t}(dx_{1:t}) \\ &= \phi_{0:t|t-1;\theta}(h_{t;\theta} f_{0:t}), \end{aligned} \tag{8}$$

where  $h_{t;\theta}(x_{0:t}) = \sum_{s=0}^{t-1} \tilde{h}_{s;\theta}(x_{s:s+1})$ , with

$$\tilde{h}_{s;\theta}(x_{s:s+1}) = \nabla_\theta \log g_{s;\theta}(x_s) + \nabla_\theta \log q_\theta(x_s, x_{s+1}), \tag{9}$$

is complete data score function. The identity (8) is established by, first, swapping, on the left hand side, the order of differentiation and integration and, second, dividing and multiplying the integrand by the complete data likelihood. In the case  $f_{0:t} = \mathbb{1}_{\mathcal{X}^{t+1}}$ , (8) coincides with Fisher’s identity (see e.g. Cappé et al. 2005, Proposition 10.1.6). Now, using (8), we obtain straightforwardly

$$\begin{aligned} \nabla_\theta \phi_{0:t|t-1;\theta} f_{0:t} &= \nabla_\theta \frac{1}{L_\theta(y_{0:t-1})} \int f_{0:t}(x_{0:t}) \chi(dx_0) \prod_{s=0}^{t-1} g_{s;\theta}(x_s) q_\theta(x_s, x_{s+1}) \mu^{\otimes t}(dx_{1:t}) \\ &= \frac{1}{L_\theta(y_{0:t-1})} \nabla_\theta \int f_{0:t}(x_{0:t}) \chi(dx_0) \prod_{s=0}^{t-1} g_{s;\theta}(x_s) q_\theta(x_s, x_{s+1}) \mu^{\otimes t}(dx_{1:t}) \\ &\quad - \phi_{0:t|t-1;\theta} f_{0:t} \frac{1}{L_\theta(y_{0:t-1})} \nabla_\theta \int \chi(dx_0) \prod_{s=0}^{t-1} g_{s;\theta}(x_s) q_\theta(x_s, x_{s+1}) \mu^{\otimes t}(dx_{1:t}) \\ &= \phi_{0:t|t-1;\theta}(h_{t;\theta} f_{0:t}) - \phi_{0:t|t-1;\theta} f_{0:t} \times \phi_{0:t|t-1;\theta} h_{t;\theta}. \end{aligned} \tag{10}$$

In other words,  $\nabla_\theta \phi_{0:t|t-1;\theta} f_{0:t}$  coincides with the covariance of  $f_{0:t}$  and the complete data score function  $h_{t;\theta}$  under  $\phi_{0:t|t-1;\theta}$ . Thus, in order to obtain, from (10), an expression of the gradient of the predictor distribution  $\pi_{t;\theta}$ , which is the marginal of  $\phi_{0:t|t-1;\theta}$  with respect to the last state, we consider  $f_t \in F(\mathcal{X})$ , let  $f_{0:t}(x_{0:t}) = f_t(x_t)$ , and use the backward decomposition (6) to obtain

$$\nabla_\theta \pi_{t;\theta} f_t = \pi_{t;\theta} \{(\mathbf{T}_{t;\theta} h_{t;\theta} - \pi_{t;\theta} \mathbf{T}_{t;\theta} h_{t;\theta}) f_t\}. \tag{11}$$

The tangent filter  $\eta_{t;\theta}$  is now defined as the signed measure

$$\eta_{t;\theta} f_t := \nabla_\theta \pi_{t;\theta} f_t, \quad f_t \in F(\mathcal{X}).$$

Since the predictors as well as the backward statistics  $\{\mathbf{T}_{t;\theta} h_{t;\theta}\}_{t \in \mathbb{N}}$  may be updated recursively through the filtering recursion (4)–(5) and the recursion (7), respectively, the identity (11) allows for recursive updating of the filter derivatives as well. Still,



since neither the predictors nor the backward statistics are generally available in a closed form, we need to resort to approximations. This will be the topic of the next section.

### 2.3 Particle-based smoothing in SSMs

Clearly, the approach to filter derivative estimation outlined above requires access to the filter and predictor distributions as well as the backward statistics. However, these quantities are available in a closed form only in very exceptional cases, e.g. when the state-space model is linear Gaussian or the state space  $X$  is a finite set. In the general case, we have to rely on approximations, and in this paper we employ SMC methods for this purpose. As mentioned in the introduction, SMC methods are genetic-type algorithms propagating recursively a random sample of particles with associated weights through repeated importance sampling.

In the following, we assume that all random variables are well defined on a common probability space  $(\Omega, \mathcal{F}, \mathbb{P})$ .

#### 2.3.1 The bootstrap particle filter

The bootstrap particle filter (Gordon et al. 1993) updates sequentially in time a set of particles with associated importance weights in order to approximate the filter and prediction distribution flows  $\{\phi_{t;\theta}\}_{t \in \mathbb{N}}$  and  $\{\pi_{t;\theta}\}_{t \in \mathbb{N}}$ , respectively, given the sequence  $\{y_t\}_{t \in \mathbb{N}}$  of observations. These methods are best illustrated in a recursive manner. Thus, assume that we have at hand a particle sample  $\{\xi_t^i\}_{i=1}^N$  approximating the predictor  $\pi_{t;\theta}$  in the sense that for all  $f \in F(\mathcal{X})$ , as  $N$  tends to infinity,

$$\pi_{t;\theta}^N f := \frac{1}{N} \sum_{i=1}^N f(\xi_t^i) \simeq \pi_{t;\theta} f.$$

Using the *updating step* (4) of the filtering recursion we can, as soon as the observation  $y_t$  becomes available, transform the uniformly weighted particle sample  $\{\xi_t^i\}_{i=1}^N$  into a weighted particle sample approximating  $\phi_{t;\theta}$  by associating each particle in the sample with the importance weight  $\omega_t^i := g_{t;\theta}(\xi_t^i)$ . Now, the weighted particle sample  $\{(\omega_t^i, \xi_t^i)\}_{i=1}^N$  targets  $\phi_{t;\theta}$  in the sense that for all  $f \in F(\mathcal{X})$ , as  $N$  tends to infinity,

$$\phi_{t;\theta}^N f := \sum_{i=1}^N \frac{\omega_t^i}{\Omega_t} f(\xi_t^i) \simeq \phi_{t;\theta} f,$$

where  $\Omega_t := \sum_{i=1}^N \omega_t^i$  denotes the weight sum. In order to form a uniformly weighted particle sample  $\{\xi_{t+1}^i\}_{i=1}^N$  targeting the next predictor  $\pi_{t+1;\theta}$ , we plug in the filter approximation  $\phi_{t;\theta}^N$  into the *prediction step* (5), which results in an approximation of  $\pi_{t+1;\theta}$  given by a mixture proportional to  $\sum_{i=1}^N \omega_t^i \mathbf{Q}_\theta(\xi_t^i, \cdot)$ . The particle cloud is then updated through sampling from this mixture, an operation that is typically carried through in two steps, *selection* and *mutation*, which are analogous to updating and

prediction. In the selection step, a set  $\{I_{t+1}^i\}_{i=1}^N$  of indices are drawn from  $\Pr(\{\omega_t^i\}_{i=1}^N)$ , where for any positive numbers  $\{a_i\}_{i=1}^k$ ,  $\Pr(\{a_i\}_{i=1}^k)$  denotes the categorical distribution on  $\{1, \dots, k\}$  induced by the probabilities proportional to  $\{a_i\}_{i=1}^k$ . After this, the mutation step propagates the particles forwards according to the dynamics of the state process, i.e. for all  $i \in \{1, \dots, N\}$ ,

$$\xi_{t+1}^i \sim \mathbf{Q}_\theta \left( \xi_t^{I_{t+1}^i}, \cdot \right).$$

The algorithm is initialised by drawing  $\{\xi_0^i\}_{i=1}^N \sim \chi^{\otimes N}$ .

The bootstrap particle filter is well suited for approximating the filter and prediction distribution flows, but it can also be used for estimating the flow of joint-smoothing distributions. In the *Poor man’s smoother*, this is done by tracing, backwards in time, the genealogical history of the particles, and using the empirical measures associated with these ancestral lineages as approximations of the joint-smoothing distributions. Unfortunately, the repeated resampling operations of bootstrap particle filter collapse these trajectories in the long run, implying, for large  $t$ , the existence of a random time point  $T < t$  before which all the trajectories coincide, i.e. all the particles  $\{\xi_t^i\}_{i=1}^N$  have a common time  $T$  ancestor. Thus, this naive approach leads to a severely depleted estimator; see Kitagawa (1996), Kitagawa and Sato (2001) for discussions of this particle path degeneracy phenomenon and Jacob et al. (2013), Olsson et al. (2008), Olsson and Westerborn (2017) for theoretical analyses of the same.

### 2.4 Online smoothing of additive functionals

A way of detouring the particle path degeneracy phenomenon goes via the backward decomposition (6). Moreover, as we will see next, (7) provides, when the objective function of interest is of additive form, a means for recursive, non-collapsed particle smoothing. This requires each backward kernel  $\overleftarrow{\mathbf{Q}}_{\phi_s; \theta}$ ,  $s \in \mathbb{N}$ , to be approximated, which is naturally done through

$$\overleftarrow{\mathbf{Q}}_{\phi_s; \theta} f(x) = \sum_{i=1}^N \frac{\omega_s^i q_\theta(\xi_s^i, x)}{\sum_{\ell=1}^N \omega_s^\ell q_\theta(\xi_s^\ell, x)} f(\xi_s^i), \quad (x, f) \in X \times F(\mathcal{X}).$$

The forward-only smoothing algorithm proposed by Del Moral et al. (2010) consists in replacing, in (7),  $\overleftarrow{\mathbf{Q}}_{\phi_t; \theta}$  by the approximation  $\overleftarrow{\mathbf{Q}}_{\phi_s; \theta}^N$ . Proceeding again recursively and assuming that we have at hand estimates  $\{\tilde{\tau}_t^i\}_{i=1}^N$  of  $\{\mathbf{T}_t h_t(\xi_t^i)\}_{i=1}^N$ , we obtain the updating formula

$$\tilde{\tau}_{t+1}^i = \sum_{j=1}^N \frac{\omega_t^j q_\theta(\xi_t^j, \xi_{t+1}^i)}{\sum_{\ell=1}^N \omega_t^\ell q_\theta(\xi_t^\ell, \xi_{t+1}^i)} \left( \tilde{\tau}_t^j + \tilde{h}_t(\xi_t^j, \xi_{t+1}^i) \right). \tag{12}$$

The recursion is initialised by setting  $\tilde{\tau}_0^i = 0$  for all  $i \in \{1, \dots, N\}$ . On the basis of these estimates, each expectation  $\phi_{0:t|t-1;\theta} h_t$  may then be approximated by

$$\phi_{0:t|t-1;\theta} h_t := \frac{1}{N} \sum_{i=1}^N \tilde{\tau}_t^i.$$

This method allows for online estimation of  $\phi_{0:t+1|t;\theta} h_{t+1}$ . It also has the appealing property that only the current statistics  $\{\tilde{\tau}_t^i\}_{i=1}^N$  and particle cloud  $\{(\xi_t^i, \omega_t^i)\}_{i=1}^N$  need to be stored. However, since the updating formula (12) requires, at each time step, a sum of  $N$  terms to be computed for each particle, the resulting algorithm has a computational complexity that grows *quadratically* with the number  $N$  of particles. Needless to say, this forces the user to keep the particle sample size relatively small, resulting in low precision.

### 2.4.1 The PaRIS

In order to increase the accuracy for a given computational budget, [Olsson and Westerborn \(2017\)](#) replace (12) by the Monte Carlo estimate

$$\tau_{t+1}^i = \frac{1}{\tilde{N}} \sum_{j=1}^{\tilde{N}} \left( \tau_t^{J^{(i,j)}} + \tilde{h}_t(\xi_t^{J^{(i,j)}}, \xi_{t+1}^i) \right), \tag{13}$$

where the sample size  $\tilde{N} \in \mathbb{N}^*$  (the so-called *precision parameter*) is typically very small compared to  $N$  and  $\{J_{t+1}^{(i,j)}\}_{j=1}^{\tilde{N}}$  are i.i.d. samples from  $\Pr(\{\omega_t^\ell q_\theta(\xi_t^\ell, \xi_{t+1}^i)\}_{\ell=1}^N)$ . Using the updated  $\{\tau_{t+1}^i\}_{i=1}^N$ , an estimate of  $\phi_{0:t+1|t;\theta} h_{t+1} = \pi_{t+1;\theta} \mathbf{T}_{t+1;\theta} h_{t+1}$  is obtained as  $N^{-1} \sum_{i=1}^N \tau_{t+1}^i$ . As before, the algorithm is initialised by setting  $\tau_0^i = 0$  for  $i \in \{1, \dots, N\}$ .

In its most basic form, also the previous approach has an  $O(N^2)$  complexity since it requires the normalising constant  $\sum_{\ell=1}^N \omega_t^\ell q_\theta(\xi_t^\ell, \xi_{t+1}^i)$  of the distribution  $\Pr(\{\omega_t^\ell q_\theta(\xi_t^\ell, \xi_{t+1}^i)\}_{\ell=1}^N)$  to be computed particle-wise, i.e. for all  $\xi_{t+1}^i$ . In order to speed up the algorithm, we apply the accept–reject sampling approach proposed by [Douc et al. \(2011\)](#). This technique presupposes that there exists a constant  $\bar{\epsilon} > 0$  such that  $q_\theta(x, x') \leq \bar{\epsilon}$  for all  $(x, x') \in \mathcal{X}^2$ , an assumption that is satisfied for most models. Then, in order to sample from  $\Pr(\{\omega_t^\ell q_\theta(\xi_t^\ell, \xi_{t+1}^i)\}_{\ell=1}^N)$ , a candidate  $J^* \sim \Pr(\{\omega_t^i\}_{i=1}^N)$  is accepted with probability  $q_\theta(\xi_t^{J^*}, \xi_{t+1}^i) / \bar{\epsilon}$ . This procedure is repeated until acceptance. Under certain mixing assumptions (see Assumption 2 below), it can be shown (see Douc et al. (2011, Proposition 2) and Olsson and Westerborn (2017, Theorem 10)) that the expected number of trials needed for this approach, which is referred to as the *particle-based, rapid incremental smoother* (PaRIS), to update  $\{\tau_t^i\}_{i=1}^N$  to  $\{\tau_{t+1}^i\}_{i=1}^N$  is  $O(\tilde{N}N)$ .

A key discovery of [Olsson and Westerborn \(2017\)](#) is that the algorithm converges, as  $N$  tends to infinity, for all fixed  $\tilde{N} \in \mathbb{N}^*$ . In particular, Olsson and Westerborn (2017, Corollary 5) provide a CLT with an asymptotic variance that may, as long as  $\tilde{N} \geq 2$ ,

be shown to grow only *linearly* with  $t$ , which is optimal for a Monte Carlo estimator on the product space (we refer to [Olsson and Westerborn 2017](#), Section 1, for a discussion). In addition, the additional asymptotic variance imposed by the “de-Rao–Blackwellisation” strategy (13) is inversely proportional to  $\tilde{N} - 1$ , which suggests that  $\tilde{N}$  should be kept at a modest value (in fact, using, say,  $\tilde{N} \in \{2, 3\}$  works typically well in simulations). On the other hand, in the case  $\tilde{N} = 1$ , the estimator suffers from a degeneracy problem that resembles closely that of the Poor man’s smoother, implying an asymptotic variance that grows quadratically fast with  $t$ . Similar stochastic stability aspects of the tangent filter estimator proposed in the present paper will be discussed in Sect. 3.2.

### 3 Tangent filter estimation: main results

#### 3.1 PaRIS-based estimation of tangent filters

The estimator that we propose is based on the identity (11), which we for a given  $\theta$  approximate online using the PaRIS algorithm. More specifically, we estimate, on-the-fly as new observations appear, the flow  $\{\eta_{t;\theta}\}_{t \in \mathbb{N}}$  of tangent filters. This is done by running a particle filter targeting the predictor distributions while updating estimates  $\{\tau_t^i\}_{i=1}^N$  of  $\{\mathbf{T}_{t;\theta} h_{t;\theta}(\xi_t^i)\}_{i=1}^N$ , where  $h_{t;\theta}$  is the complete data score given by (9). The statistics  $\{\tau_t^i\}_{i=1}^N$  are updated efficiently using the PaRIS updating rule (13). At each time-step  $t$ , our algorithm returns

$$\eta_{t;\theta}^N f_t = \frac{1}{N} \sum_{i=1}^N \left( \tau_t^i - \frac{1}{N} \sum_{j=1}^N \tau_t^j \right) f_t(\xi_t^i)$$

as an estimate of  $\eta_{t;\theta} f_t$  for any  $f_t \in \mathbf{F}(\mathcal{X})$ . We summarise the algorithm in the pseudo-code below, where every line with  $i$  should be executed for all  $i \in \{1, \dots, N\}$ .

- 1: draw  $\xi_0^i \sim \chi$
- 2: set  $\tau_0^i \leftarrow 0$
- 3: **for**  $t \leftarrow 0, 1, 2, \dots$  **do**
- 4:   set  $\omega_t^i \leftarrow g_{t;\theta}(\xi_t^i)$
- 5:   draw  $I_{t+1}^i \sim \Pr\left(\{\omega_t^\ell\}_{\ell=1}^N\right)$
- 6:   draw  $\xi_{t+1}^i \sim q_\theta\left(\xi_t^{I_{t+1}^i}, \cdot\right)$
- 7:   **for**  $j \leftarrow 1, \dots, \tilde{N}$  **do**
- 8:     draw  $J_{t+1}^{(i,j)} \sim \Pr\left(\{\omega_t^\ell q_\theta(\xi_t^\ell, \xi_{t+1}^i)\}_{\ell=1}^N\right)$
- 9:   **end for**
- 10:   set  $\tau_{t+1}^i \leftarrow \frac{1}{N} \sum_{j=1}^{\tilde{N}} \left( \tau_t^{J_{t+1}^{(i,j)}} + \tilde{h}_t\left(\xi_t^{J_{t+1}^{(i,j)}}, \xi_{t+1}^i\right) \right)$
- 11:   set  $\bar{\tau}_{t+1} \leftarrow \frac{1}{N} \sum_{\ell=1}^N \tau_{t+1}^\ell$
- 12:   set  $\eta_{t+1;\theta}^N \leftarrow \frac{1}{N} \sum_{\ell=1}^N (\tau_{t+1}^\ell - \bar{\tau}_{t+1}) \delta_{\xi_{t+1}^\ell}$

13: **end for**

**Remark 1** The previous algorithm can be viewed as a “de-Rao–Blackwellisation” of the tangent filter estimator proposed in [Del Moral et al. \(2015\)](#), which is based on the updating rule (12) instead of (13). With the exception of this important difference, the two algorithms follow the same recursion. Consequently, the algorithm of [Del Moral et al. \(2015\)](#) has an  $O(N^2)$  complexity.

### 3.2 Theoretical results

The first part of the convergence analysis of our algorithm will be carried through under the following, relatively mild, assumption.

- Assumption 1** (i) For all  $t \in \mathbb{N}$  and  $\theta \in \Theta$ ,  $g_{t;\theta}$  is a positive and bounded measurable function.  
 (ii) For all  $\theta \in \Theta$ ,  $q_\theta \in F(\mathcal{X}^{\otimes 2})$ .

Assumption 1(i) implies finiteness and positiveness of the particle importance weights, while Assumption 1(ii) allows, apart from some technical arguments needed in the proofs, the accept–reject sampling technique mentioned in Sect. 2.4 to be used.

We begin by establishing an exponential concentration inequality, valid for all finite particle sample sizes  $N$ , for our PaRIS-based tangent filter estimator. This yields the strong ( $\mathbb{P}$ -a.s.) consistency (as  $N$  tends to infinity) of the algorithm as a corollary. In addition, we state and establish a CLT with an asymptotic variance given by a sum of the asymptotic variance of the estimator proposed by [Del Moral et al. \(2015\)](#) and one term reflecting the additional variance introduced by the supplementary sampling performed in the PaRIS-version. Finally, we are able to derive, by operating with strong mixing assumptions, a time-uniform bound on the asymptotic variance, which is inversely proportional to  $\tilde{N} - 1$ . This guarantees that the algorithm is stochastically stable in the long run. All proofs, which use results obtained by [Olsson and Westerborn \(2017\)](#), are postponed to “Appendix A”.

#### 3.2.1 Results on convergence

**Theorem 1** *Let Assumption 1 hold. Then, for all  $t \in \mathbb{N}$ ,  $\theta \in \Theta$ ,  $f_t \in F(\mathcal{X})$ , and  $\tilde{N} \in \mathbb{N}^*$  there exists  $(c_t, \tilde{c}_t) \in (\mathbb{R}_+^*)^2$  (depending on  $\theta$ ,  $\tilde{N}$ , and  $f_t$ ) such that for all  $N \in \mathbb{N}^*$  and  $\varepsilon \in \mathbb{R}_+$ ,*

$$\mathbb{P} \left( \left| \eta_{t;\theta}^N f_t - \eta_{t;\theta} f_t \right| \geq \varepsilon \right) \leq c_t \exp(-\tilde{c}_t N \varepsilon (\varepsilon \wedge 1)).$$

The strong consistency of the estimator follows.

**Corollary 1** *Let Assumption 1 hold. Then, for all  $t \in \mathbb{N}$ ,  $\theta \in \Theta$ ,  $f_t \in F(\mathcal{X})$ , and  $\tilde{N} \in \mathbb{N}^*$  it holds,  $\mathbb{P}$ -a.s.,*

$$\lim_{N \rightarrow \infty} \eta_{t;\theta}^N f_t = \eta_{t;\theta} f_t.$$

We now set focus on weak convergence. For this purpose, we introduce, for all  $t \in \mathbb{N}$ , the unnormalised kernels

$$\mathbf{L}_{t;\theta} f(x) := g_{t;\theta}(x) \mathbf{Q}_\theta f(x), \quad (x, f) \in X \times F(\mathcal{X}),$$

and, for all  $(s, t) \in \mathbb{N}^2$  such that  $s \leq t$ , the retro-prospective kernels

$$\begin{aligned} \mathbf{D}_{s,t;\theta} f(x_s) &:= \int \int f(x_{0:t}) g_{t;\theta}(x_t) \mathbf{T}_{s;\theta}(x_s, dx_{0:s-1}) \mathbf{L}_{s;\theta} \otimes \cdots \otimes \mathbf{L}_{t-1;\theta}(x_s, dx_{s+1:t}), \\ \tilde{\mathbf{D}}_{s,t;\theta} f(x_s) &:= \mathbf{D}_{s,t;\theta}(f - \phi_{0:t|t-1;\theta} f)(x_s), \end{aligned}$$

for  $(x_s, f) \in X \times F(\mathcal{X}^{\otimes(t+1)})$ , operating simultaneously in the backward and forward directions.

**Theorem 2** *Let Assumption 1 hold. For all  $t \in \mathbb{N}$ ,  $\theta \in \Theta$ ,  $\tilde{N} \in \mathbb{N}^*$ , and  $f_t \in F(\mathcal{X})$ , as  $N \rightarrow \infty$ ,*

$$\sqrt{N} \left( \eta_{t;\theta}^N f_t - \eta_{t;\theta} f_t \right) \xrightarrow{\mathcal{D}} \sigma_{t;\theta}(f_t) Z,$$

where  $Z$  is a standard Gaussian random variable and

$$\sigma_{t;\theta}^2(f_t) := \tilde{\sigma}_{t;\theta}^2(f_t) + \sum_{s=0}^{t-1} \sum_{\ell=0}^s \tilde{N}^{\ell-(s+1)} \varsigma_{s,\ell,t;\theta}(f_t), \tag{14}$$

where

$$\tilde{\sigma}_{t;\theta}^2(f_t) := \sum_{s=0}^{t-1} \frac{\pi_{s+1;\theta} \tilde{\mathbf{D}}_{s+1,t;\theta}^2 \{ (h_{t;\theta} - \pi_{t;\theta} \mathbf{T}_{t;\theta} h_{t;\theta})(f_t - \pi_{t;\theta} f_t) \}}{(\pi_{s+1;\theta} \mathbf{L}_{s+1;\theta} \cdots \mathbf{L}_{t-1;\theta} \mathbb{1}_X)^2} \tag{15}$$

is the asymptotic variance of the tangent filter estimator proposed by [Del Moral et al. \(2015\)](#) and

$$\begin{aligned} \varsigma_{s,\ell,t;\theta}(f_t) &:= \\ & \frac{\pi_{\ell+1;\theta} \{ \tilde{\mathbf{Q}}_{\phi_{\ell;\theta}}(\mathbf{T}_{\ell;\theta} h_{\ell;\theta} + \tilde{h}_{\ell;\theta} - \mathbf{T}_{\ell+1;\theta} h_{\ell+1;\theta})^2 \mathbf{L}_{\ell+1;\theta} \cdots \mathbf{L}_{s;\theta} (\mathbf{L}_{s+1;\theta} \cdots \mathbf{L}_{t-1;\theta} \{ f_t - \pi_{t;\theta} f_t \})^2 \}}{(\pi_{\ell+1;\theta} \mathbf{L}_{\ell+1;\theta} \cdots \mathbf{L}_{s;\theta} \mathbb{1}_X) (\pi_{s+1;\theta} \mathbf{L}_{s+1;\theta} \cdots \mathbf{L}_{t-1;\theta} \mathbb{1}_X)^2}. \end{aligned}$$

**Remark 2** We note that for all  $t \in \mathbb{N}$ ,

$$\lim_{\tilde{N} \rightarrow \infty} \sigma_{t;\theta}^2(f_t) = \tilde{\sigma}_{t;\theta}^2(f_t).$$

This is in line with our expectations, as the tangent filter estimator proposed by [Del Moral et al. \(2015\)](#) can be viewed as a Rao–Blackwellisation of our estimator.

### 3.2.2 Results on stochastic stability

We now set focus on the numerical stability of the algorithm and show that the error of the tangent filter estimator proposed by us stays uniformly bounded in time. There are several ways of formulating such stability, and in the present paper, we follow [Del Moral and Guionnet \(2001\)](#), [Douc et al. \(2014\)](#) and establish a time-uniform bound on the asymptotic variance of the estimator. The analysis requires Assumption 1 to be sharpened as follows.

**Assumption 2** (i) There exists  $\rho > 1$  such that for all  $\theta \in \Theta$  and  $(x, \tilde{x}) \in X^2$ ,

$$\rho^{-1} \leq q_\theta(x, x') \leq \rho.$$

(ii) There exist constants  $0 < \underline{\delta} < \bar{\delta} < \infty$  such that for all  $t \in \mathbb{N}$  all  $\theta \in \Theta$ , and  $x \in X$ ,

$$\underline{\delta} \leq g_{t;\theta}(x) \leq \bar{\delta}.$$

Under Assumption 2 we define

$$\varrho := 1 - \rho^{-2}.$$

These—admittedly very restrictive—*strong mixing assumptions* require typically the state space  $X$  of the underlying Markov chain to be a compact set. Nevertheless, such assumptions are standard in the literature of SMC analysis; see [Del Moral and Guionnet \(2001\)](#) and, e.g. [Crisan and Heine \(2008\)](#), [Douc et al. \(2014\)](#) for refinements. Finally, we assume that the terms (9) of the complete data score function are uniformly bounded.

**Assumption 3** There exists a constant  $|\tilde{h}|_\infty \in \mathbb{R}_+$  such that for all  $s \in \mathbb{N}$  and  $\theta \in \Theta$ ,

$$\text{osc}(\tilde{h}_{s;\theta}) \leq |\tilde{h}|_\infty.$$

In conformity with the strong mixing assumptions, also Assumption 3 requires typically the state space  $X$  to be compact.

The previous assumptions allow us to establish the following theorem, providing a time-uniform bound on the sequence  $\{\sigma_{t;\theta}^2(f)\}_{t \in \mathbb{N}}$  of asymptotic variances. Generally, when the objective function is of additive form, the variance of a Monte Carlo estimator increases typically linearly with the number of terms, or, the dimension of the underlying product space (see e.g. [Olsson and Westerborn 2017](#), Section 1, for a discussion). Thus, as our estimator is based on the identities (10) and (11), which express the tangent filter in terms of additive smoothed expectations over the path space, it may seem surprising that we are able to derive a time-uniform bound on the asymptotic variance. However, recall from (10) that the tangent filter  $\eta_{t;\theta} f_t$  coincides with the covariance of  $f_t$  and the complete data score function  $h_{t;\theta}$  under  $\phi_{0:t|t-1;\theta}$ . Thus, when the model exhibits forgetting properties—which is the case under the strong mixing assumptions

above—the sequence  $\{\eta_{t;\theta} f_t\}_{t \in \mathbb{N}}$  stays bounded, and time-uniform convergence of the Monte Carlo estimator is hence within reach.

**Theorem 3** *Let Assumptions 2 and 3 hold. Then there exist constants  $(c, \tilde{c}) \in \mathbb{R}_+^2$  such that for all  $t \in \mathbb{N}$ ,  $\theta \in \Theta$ ,  $\tilde{N} \geq 2$ , and  $f_t \in F(\mathcal{X})$ ,*

$$\sigma_{t;\theta}^2(f_t) \leq \|f_t\|_\infty^2 \left( c + \tilde{c} \frac{1}{(\tilde{N} - 1)(1 - \rho^2)} \right)$$

and

$$\tilde{\sigma}_{t;\theta}^2(f_t) \leq c \|f_t\|_\infty^2. \quad (16)$$

The bound (16) on the first term (15) of the asymptotic variance, corresponding to the variance of the Rao–Blackwellised algorithm, was derived by [Del Moral et al. \(2015\)](#). As expected from the theoretical results obtained by [Olsson and Westerborn \(2017\)](#), the bound on the second term is inversely proportional to  $\tilde{N} - 1$ , which indicates that boosting the number  $\tilde{N}$  of PaRISian backward draws from a moderate to a very large value will have negligible effect on the precision. We hence recommend keeping  $\tilde{N}$  low, say lower than 5, in order to gain computational efficiency. A more detailed discussion on the choice of  $\tilde{N}$  can be found in [Olsson and Westerborn \(2017\)](#).

## 4 Application to recursive maximum likelihood estimation

### 4.1 Batch mode implementation

Given a fixed data–record  $y_{0:t}$ , our goal is to perform maximum likelihood estimation of  $\theta$ , i.e. to find the vector of parameters  $\theta^* \in \Theta$  such that  $\theta^* = \arg \max_{\theta \in \Theta} L_\theta(y_{0:t})$ , where  $L_\theta(y_{0:t})$  is the likelihood of the observed data given in (3). Equivalently, we may maximise the log-likelihood  $\ell_\theta(y_{0:t}) = \log L_\theta(y_{0:t})$ . There are many different approaches to such maximisation; see ([Cappé et al. 2005](#), Chapter 10) for a more general overview of different approaches.

Here, we will focus on the following *Robbins–Monro scheme*: at iteration  $n$ , let

$$\theta_n = \theta_{n-1} + \gamma_n Z_n,$$

where  $Z_n$  is a noisy measurement of  $\nabla_\theta \ell_\theta(y_{0:t})|_{\theta=\theta_{n-1}}$ , i.e. the score of the observed data evaluated in  $\theta = \theta_{n-1}$ , and  $\{\gamma_n\}_{n \in \mathbb{N}^*}$  a sequence of positive step-sizes satisfying the regular stochastic approximation requirements  $\sum_{n=1}^{\infty} \gamma_n = \infty$  and  $\sum_{n=1}^{\infty} \gamma_n^2 < \infty$ . Note that this approach requires approximation of  $Z_n$  at each iteration of the algorithm. If the number of observations is very large computing  $Z_n$  is costly, and since many iterations are often required for convergence, this results in an impractical algorithm. Moreover, if we receive a new observation,  $Z_n$  needs to be recalculated which turns the procedure into an offline algorithm.



### 4.2 Online implementation: PaRISian RML

We sketch the basic principles of recursive maximum likelihood. First note that since  $\{(X_t, Y_t)\}_{t \in \mathbb{N}}$  is a Markov chain, the quadruple  $\{(X_t, Y_t, \pi_t; \theta, \eta_t; \theta)\}_{t \in \mathbb{N}}$  forms, by (4) and (5) and for all  $\theta \in \Theta$ , a Markov chain as well. In the case where the state space  $X$  is a finite set, Le Gland and Mevel (1996) showed that this chain is ergodic under certain mixing assumptions. This result was later extended by Douc and Matias (2001) to the case where  $X$  is compact. Now, assume that data is generated by a model parameterised by a distinctive  $\theta_* \in \Theta$ ; then, denoting by  $\Pi_{\theta, \theta_*}$  the stationary distribution of this chain and  $\tilde{\Pi}_{\theta, \theta_*}(\cdot) = \Pi_{\theta, \theta_*}(X \times \cdot)$  its marginal with respect to the last three components, it holds for all  $\theta \in \Theta$ ,  $\mathbb{P}$ -a.s.,

$$\begin{aligned} \lim_{t \rightarrow \infty} \frac{1}{t} \nabla_{\theta} \ell_{\theta}(Y_{0:t}) &= \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{s=0}^t \nabla_{\theta} \ell_{\theta}(Y_s \mid Y_{0:s-1}) \\ &= \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{s=0}^t \frac{\pi_{s;\theta}(\nabla_{\theta} g_{s;\theta}) + \eta_{s;\theta} g_{s;\theta}}{\pi_{s;\theta} g_{s;\theta}} \\ &= \iiint \frac{\pi(\nabla_{\theta} g_{\theta}(\cdot, y)) + \eta(g_{\theta}(\cdot, y))}{\pi(g_{\theta}(\cdot, y))} \tilde{\Pi}_{\theta, \theta_*}(d(y, \pi, \eta)) =: \lambda(\theta, \theta_*), \end{aligned}$$

where  $\pi_{s;\theta} g_{s;\theta}$ ,  $\eta_{s;\theta} g_{s;\theta}$ , and  $\pi_{s;\theta} \nabla_{\theta} g_{s;\theta}$  depend implicitly on  $Y_{0:s}$ . These equations follow, in order, by the decomposition (1) of the log-likelihood, by applying the nabla operator, and by applying Birkhoff’s ergodic theorem. By identifiability (see e.g. Cappé et al. 2005, Section 12.4), the true parameter  $\theta_*$  solves  $\lambda(\theta, \theta_*) = 0$ , and the task of solving this equation may be cast into the framework of *stochastic approximation* and the *Robbins–Monro algorithm*

$$\theta_{t+1} = \theta_t + \gamma_{t+1} \zeta_{t+1}, \quad t \in \mathbb{N}, \tag{17}$$

where  $\zeta_{t+1}$  is a noisy observation of  $\lambda(\theta_t, \theta_*)$ . Ideally, such a noisy observation would be formed by estimating  $\lambda(\theta_t, \theta_*)$  by a draw from  $\tilde{\Pi}_{\theta_t, \theta_*}$ , which is, needless to say, not possible in practice. Thus, by introducing one more level of approximation and allowing for Markovian perturbations, one may instead estimate, following (11),  $\lambda(\theta_t, \theta_*)$  by

$$\zeta_{t+1} := \frac{\zeta_{t+1}^1 + \zeta_{t+1}^2}{\zeta_{t+1}^3},$$

where

$$\begin{aligned} \zeta_{t+1}^1 &:= \pi_{t+1}(\nabla_{\theta} g_{t+1;\theta} \mid \theta = \theta_t), \\ \zeta_{t+1}^2 &:= \eta_{t+1} g_{t+1;\theta_t}, \\ \zeta_{t+1}^3 &:= \pi_{t+1} g_{t+1;\theta_t}, \end{aligned}$$

with

$$\eta_{t+1} := \pi_t \{(T_t - \pi_t T_t) g_{t;\theta_t}\},$$

and statistics  $\{T_t\}_{t \in \mathbb{N}}$  being, following (7), updated recursively according to

$$T_{t+1}(x_{t+1}) = \int \{T_t(x_t) + \tilde{h}_{t;\theta_{t+1}}(x_{t:t+1})\} \overleftarrow{\mathbf{Q}}_{\phi_t}(x_{t+1}, dx_t), \quad x_{t+1} \in \mathcal{X}, \quad (18)$$

and initialised as  $T_0 \equiv 0$ . The statistic (18) depends on the new observation  $Y_{t+1}$  through the term  $\tilde{h}_{t;\theta_{t+1}}$ . In addition, the measures  $\{\phi_t\}_{t \in \mathbb{N}}$  and  $\{\pi_t\}_{t \in \mathbb{N}}$  are updated according to the recursion

$$\begin{aligned} \phi_t f &= \frac{\pi_t(g_{t;\theta_t} f)}{\pi_t g_{t;\theta_t}}, \quad f \in F(\mathcal{X}), \\ \pi_{t+1} f &= \phi_t \mathbf{Q}_{\theta_t} f, \end{aligned} \quad (19)$$

with, by convention,  $\pi_0 := \chi$ . Note that the recursion (19) updates the filter and predictor distributions on the basis of the current parameter fit. This means that for all  $t \in \mathbb{N}$ , the measures  $\phi_t$ ,  $\pi_{t+1}$  and  $\eta_{t+1}$  depend on *all* the past parameter values  $\{\theta_s\}_{s=0}^t$ .

This approach yields an online algorithm where a new observation can be incorporated into the estimator without the need of having to recalculate the latter from scratch. In the case where  $\mathcal{X}$  is finite, this algorithm was studied in [Le Gland and Mevel \(1997\)](#), which also provides assumptions under which the sequence  $\{\theta_t\}_{t \in \mathbb{N}^*}$  converges towards the parameter  $\theta_*$  (see also [Tadic 2010](#), for refinements).

In the general case, the recursions (18) and (19) are intractable. The translation of the same into the language of particles is, however, immediate by approximating (19) by a particle filter, updated according to time-varying parameters, and (18) by means of a PaRISian updating step (13), the latter executed for some suitable precision parameter  $\tilde{N} \geq 2$ .

The algorithm is initialised by some arbitrary parameter guess  $\theta_0$ . After this, the parameter fit is updated recursively as each new observation  $y_{t+1}$  is received, by first calculating a particle approximation  $\hat{\zeta}_{t+1}$  of  $\zeta_{t+1}$  and then updating the parameter according to

$$\theta_{t+1} = \theta_t + \gamma_{t+1} \hat{\zeta}_{t+1},$$

where  $\{\gamma_t\}_{t \in \mathbb{N}^*}$  satisfies the standard stochastic optimisation requirements (i.e. infinite sum and finite sum of squares) above. The algorithm, which is illustrated numerically in the next section, is detailed in pseudo-code below.

- 1: set arbitrarily  $\theta_0$
- 2: draw  $\xi_0^i \sim \chi$
- 3: set  $\tau_0^i \leftarrow 0$
- 4: **for**  $t \leftarrow 0, 1, 2, \dots$  **do**
- 5:   set  $\omega_t^i \leftarrow g_{\theta_t}(\xi_t^i, y_t)$
- 6:   draw  $I_t \sim \Pr(\{\omega_t^\ell\}_{\ell=1}^N)$
- 7:   draw  $\xi_{t+1}^i \sim q_{\theta_t}(\xi_t^{I_t}, \cdot)$
- 8:   **for**  $j \leftarrow 1, \dots, \tilde{N}$  **do**
- 9:     draw  $J_{t+1}^{(i,j)} \sim \Pr(\{\omega_t^\ell q_{\theta_t}(\xi_t^\ell, \xi_{t+1}^i)\}_{\ell=1}^N)$
- 10:   **end for**

- 11: set  $\tau_{t+1}^i \leftarrow \tilde{N}^{-1} \sum_{j=1}^{\tilde{N}} \left( \tau_{t+1}^{j(i,j)} + \tilde{h}_{t;\theta_t} \left( \xi_{t+1}^{j(i,j)}, \xi_{t+1}^i \right) \right)$
- 12: set  $\bar{\tau}_{t+1} \leftarrow N^{-1} \sum_{\ell=1}^N \tau_{t+1}^\ell$
- 13: set  $\hat{\zeta}_{t+1}^1 \leftarrow N^{-1} \sum_{\ell=1}^N \nabla g_{\theta_t} \left( \xi_{t+1}^\ell, y_{t+1} \right)$
- 14: set  $\hat{\zeta}_{t+1}^2 \leftarrow N^{-1} \sum_{\ell=1}^N \left( \tau_{t+1}^\ell - \bar{\tau}_{t+1} \right) g_{\theta_t} \left( \xi_{t+1}^\ell, y_{t+1} \right)$
- 15: set  $\hat{\zeta}_{t+1}^3 \leftarrow N^{-1} \sum_{\ell=1}^N g_{\theta_t} \left( \xi_{t+1}^\ell, y_{t+1} \right)$
- 16: set  $\theta_{t+1} \leftarrow \theta_t + \gamma_{t+1} \frac{\hat{\zeta}_{t+1}^1 + \hat{\zeta}_{t+1}^2}{\hat{\zeta}_{t+1}^3}$
- 17: **end for**

### 4.3 Simulations

We benchmark our algorithm on two different models, namely

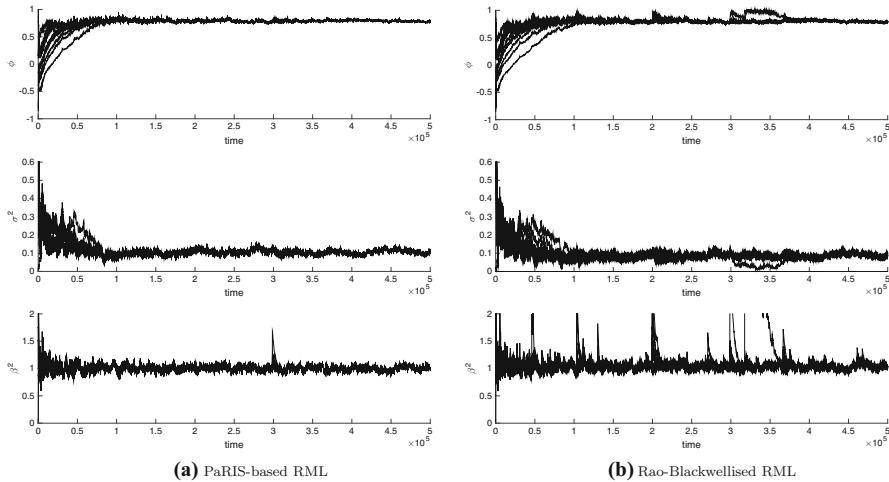
- the stochastic volatility model of [Hull and White \(1987\)](#), where we compare our estimator to that of [Del Moral et al. \(2015\)](#), and
- the simultaneous localisation and mapping (SLAM) model where we apply our algorithm as a proof of concept.

#### 4.3.1 Stochastic volatility model

Consider the stochastic volatility model

$$\begin{aligned} X_{t+1} &= \phi X_t + \sigma V_{t+1}, \\ Y_t &= \beta \exp(X_t/2) U_t, \end{aligned} \quad t \in \mathbb{N},$$

proposed by [Hull and White \(1987\)](#), where  $\{V_t\}_{t \in \mathbb{N}^*}$  and  $\{U_t\}_{t \in \mathbb{N}}$  are independent sequences of mutually independent standard Gaussian noise variables. In this model, the parameters to be estimated are  $\theta = (\phi, \sigma^2, \beta^2)$ . We compared our PaRISian RML estimator to the estimator proposed by [Del Moral et al. \(2015\)](#)—the latter being referred to as the ‘‘Rao–Blackwellised estimator’’. Since there is a significant difference in computational complexity between the algorithms, we design the particle sample size in each algorithm in such a way that the running times of the two algorithms are close to identical. With our implementation, using  $N = 100$  in the algorithm of [Del Moral et al. \(2015\)](#) yields close to the same CPU time as using  $(N, N) = (1400, 2)$  in the PaRISian RML estimator. For both algorithms, we let  $\{\gamma_t\}_{t \in \mathbb{N}} = \{t^{-0.6}\}_{t \in \mathbb{N}}$ . The algorithms were executed on an observation data record comprising 500,000 observations generated under the parameters  $\theta^* = (0.8, 0.1, 1)$ . The algorithms were run 12 times each on the same data set with the same randomised starting parameters. [Figure 1](#) displays the resulting learning trajectories. It can be seen clearly that both methods converge to the true parameter values; however, the PaRISian RML estimates exhibit significantly less variance than those produced by the Rao–Blackwellised estimator (using a considerably smaller particle sample size). Indeed, computing the sample variances over the final parameter estimates indicates an improvement by almost an order of magnitude to the benefit of the PaRISian version for the same computational



**Fig. 1** Particle learning trajectories produced by the PaRISian RML estimator (left panel) and the Rao–Blackwellised RML estimator (right panel) proposed in Del Moral et al. (2015) for, from top to bottom,  $\phi$ ,  $\sigma^2$ , and  $\beta^2$ . The former and latter algorithms used  $N = 1,400$  and  $N = 400$  particles, respectively (leading to comparable CPU times). For each algorithm, 12 replicates were generated on the same data set with different, randomised initial parameters (being the same for both algorithms). For the Rao–Blackwellised version, the plot of  $\beta^2$  does not contain the full trajectories due to very high peaks. **a** PaRIS-based RML, **b** Rao–Blackwellised RML

time:  $(.069, .181, .095) \times 10^{-4}$  and  $(.054, .164, .063) \times 10^{-3}$  for the PaRISian and Rao–Blackwellised versions, respectively. It can also be seen in Fig. 1 that the trajectories exhibit some jumps, mainly in the  $\beta^2$  variable. This occurs when the estimate of  $\zeta_{t+1}^3$  is close to zero, which corresponds to time steps when the particles fail to cover the support of the emission density. Since the computationally more efficient PaRISian version allows considerably more particles to be used for the given budget, the problem is much smaller for this algorithm; indeed, whereas the learning curves of the Rao–Blackwellised version exhibit a high degree of volatility and several large jumps (the peaks are cut for visibility) with subsequent very long excursions out in the parameter space, the PaRISian version exhibits only one significant jump of moderate size.

### 4.3.2 Simultaneous localisation and mapping (SLAM)

The simultaneous localisation and mapping (SLAM) problem is fundamental in robotics. In the version considered here, we let the state space consist of the coordinates and the bearing of a robot moving in the plane, i.e.  $\mathbf{X} = \mathbb{R}^2 \times (-\pi, \pi]$ . The prior motion of the robot is modelled as a Markov chain  $X_t := (X_t^1, X_t^2, X_t^3)$ ,  $t \in \mathbb{N}$ , on this space, defined through the state equations

$$\begin{aligned} X_{t+1}^1 &= X_t^1 + d_{t+1} \cos(X_t^3) + \epsilon_{t+1}^1, \\ X_{t+1}^2 &= X_t^2 + d_{t+1} \sin(X_t^3) + \epsilon_{t+1}^2, \end{aligned}$$

$$X_{t+1}^3 = X_t^3 + \alpha_{t+1} + \epsilon_{t+1}^3,$$

where  $\{\epsilon_t^i\}_{t \in \mathbb{N}^*}$ ,  $i \in \{1, 2, 3\}$  are independent sequences of mutually independent noise variables with known variances  $\sigma_i^2$ ,  $i \in \{1, 2, 3\}$ . The sequence  $\{(d_t, \alpha_t)\}_{t \in \mathbb{N}^*}$  provides the commands—in terms of speed and bearing changes—of the robot at each time point.

The robot observes a landmark, defined by its positions in the plane, by measuring the distance and the relative bearing to the same. Assuming  $L$  landmarks, the observations at time  $t$  are given by  $Y_t = \{Y_t^i\}_{i \in O_t}$ , where  $O_t \subseteq \{1, \dots, L\}$  is the set of observed landmarks. For each observed landmark,


$$Y_t^i = \left( \begin{array}{l} \left\| (\theta_1^i, \theta_2^i) - (X_t^1, X_t^2) \right\| + \epsilon_t^{i,1} \\ \arctan \left( \frac{\theta_2^i - X_t^2}{\theta_1^i - X_t^1} \right) - X_t^3 + \epsilon_t^{i,2} \end{array} \right),$$

where  $\|\cdot\|$  denotes the Euclidean norm,  $(\theta_1^i, \theta_2^i)$  is the location of landmark  $i$ ,  $(\epsilon_t^{i,1}, \epsilon_t^{i,2})$  are independent noise variables with known variances  $(\varsigma_1^2, \varsigma_2^2)$ , the noises of different time steps and different landmarks being independent.

In this setting, we wish to estimate the locations of all the landmarks, which implies  $2L$  unknown parameters. Note that the noise parameters in the model are assumed to be calibrated beforehand. Several existing works apply particle methods to the SLAM problem; see e.g. the *FastSLAM* (Montemerlo et al. 2002) and *FastSLAM 2.0* (Montemerlo et al. 2003) algorithms. More recently an online EM version was proposed by Le Corff et al. (2011) and the *Marginal-SLAM algorithm* (Martinez-Cantin et al. 2007) is based on an updating step similar to (17).

Using the model above, data were simulated for  $L = 9$  landmarks located according to the green dots in Fig. 2. The observations were generated by simulating 100,000 observations of the robot moving around a closed loop with a time resolution of .1 s. The noise parameters were set to  $\sigma_1 = \sigma_2 = .25$ ,  $\sigma_3 = \frac{3\pi}{180}$ ,  $\varsigma_1 = .25$ , and  $\varsigma_2 = \frac{\pi}{180}$ . The robot observes landmarks within a 30 m radius and a 180° field of vision.

Using this simulated data, we performed 20 independent runs of our PaRISian RML algorithm in order to estimate the landmark positions. For each landmark, the initial estimate of its position was drawn randomly according to a bivariate Gaussian distribution with the true landmark position as mean and the identity matrix as covariance matrix. The algorithm used  $(N, \tilde{N}) = (500, 2)$  and the particles were initialised in  $(0, 0)$  with zero bearing, i.e. the same as the robot’s starting position. Since the problem is invariant under translations and rotations, we fix two landmarks and assume that these are known a priori. The results are displayed in Fig. 2, where the estimates cluster closely around the true parameter values. The figure also shows the estimated filtered trajectory of the robot during its first 3 laps. In order to illustrate the convergence of the landmark position estimates, we consider the evolution of their MSE over time. The results are reported in Fig. 3, where the average MSE of 4 landmarks is computed based on the 20 independent runs of the algorithm. As clear from the picture, the processing of the full observation record yields MSE values that fluctuate stably around zero.

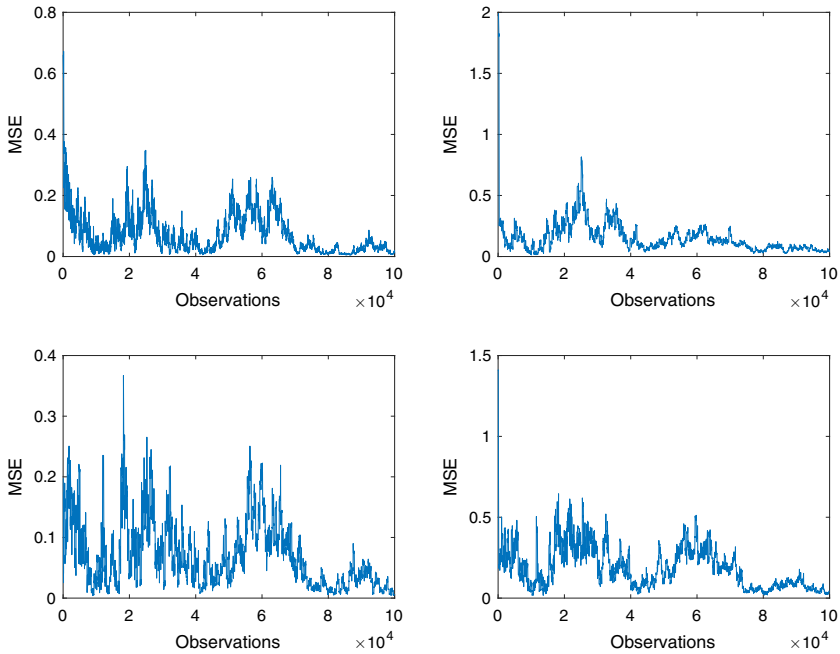


**Fig. 2** Resulting estimates of the 9 landmarks for the SLAM problem using the PaRISian RML algorithm. The circles are the true positions of the landmarks and green stars are the resulting estimates. The black dots denote the particle estimates of the robot's positions during the first three laps. Rerunning the algorithm yields similar trajectories

Needless to say, even though this prototypical solution to the SLAM seems promising, a lot of additional work is needed for obtaining an algorithm running in real applications. Moreover, the comprehensive task of benchmarking our approach against full palette of existing algorithms treating this fundamental problem falls outside the scope of the present paper.

## 5 Conclusion

We have presented a novel algorithm for efficient particle-based estimation of tangent filters in general SSMs. The algorithm involves the PaRIS algorithm proposed by [Olsson and Westerborn \(2017\)](#) as a key ingredient. The estimator is furnished with several convergence results, the main result being a CLT at the rate  $\sqrt{N}$ . The convergence analysis is driven by results obtained by [Olsson and Westerborn \(2017\)](#), [Del Moral et al. \(2015\)](#), and our time-uniform bound on the asymptotic variance extends the existing results in [Del Moral et al. \(2015\)](#). Importantly, under weak assumptions, the computational complexity of our algorithm grows only linearly with the number of particles, whereas the complexity of existing methods such as the estimator of [Del Moral et al. \(2015\)](#), which may be viewed as a Rao–Blackwellisation of our estimator, is typically quadratic in the number of particles. Thus, we may expect the computational benefit



**Fig. 3** Resulting mean square error estimate for 4 of the landmarks based on the 20 independent runs of the PaRISian RML algorithm for the SLAM problem

of the “de-Rao–Blackwellisation” that characterises PaRIS to grossly exceed the price of the additional variance introduced by the PaRISian backward sampling procedure.

In the second part of the paper, we cast our PaRISian tangent filter estimator into the framework of RML, yielding a computationally efficient and easily implemented online parameter estimation procedure. As clear from the simulations, the fact that our online estimator allows, compared to existing methods, more particles to be used for a given computational budget is of importance for the stability of the online estimates.

Appealingly, the asymptotic variance of our estimator can be bounded uniformly in time; in other words, the estimator stays stochastically stable in the long term. Needless to say, the strong mixing assumptions driving the stability analysis are restrictive, and to relax these assumptions is a natural direction of research. [Jasra \(2015\)](#) provides an  $O(t)$  bound on the asymptotic variance of the *forward-filtering backward-smoothing* (FFBSm) algorithm ([Doucet et al. 2000](#)), which is equivalent with the Rao–Blackwellised PaRIS in the case of additive state functionals (the online formulation was however found by [Del Moral et al. \(2010\)](#)), under assumptions that point to applications on non-compact state spaces, and the same approach may be applicable to our PaRISian tangent filter estimator.

Finally, it still remains to prove, under general, verifiable assumptions, the convergence of the parameter estimates produced by our PaRISian RML algorithm. Such an analysis is expected to be technically very involved, and is hence beyond the scope of the present paper. In [Le Gland and Mevel \(1997\)](#), in which the convergence of

exact RML in HMMs is established for the case of finite state space, the proof consists in showing that the extended Markov chain comprising the state and observation processes as well as the prediction and tangent filter flows is ergodic and applying standard stochastic approximation results [Delyon \(1996\)](#). A route to the convergence of our proposed PaRISian RML could be to include also the particle cloud and the particle-based statistics  $\{\tau_t^i\}_{i=1}^N$  into the Markov chain and study the ergodicity of the same. An alternative and more direct approach would be to cast the problem into the framework of *biased stochastic gradient search*, which was analysed recently by [Tadić and Doucet \(2017\)](#).

### A Proofs

Define for all  $t \in \mathbb{N}$  and  $\theta \in \Theta$ ,

$$\mathbf{L}_{t;\theta} : X \times \mathcal{X} \ni (x, A) \mapsto g_{t;\theta}(x)\mathbf{Q}_\theta(x, A).$$

(Note that our definition of  $\mathbf{L}_t$  differs from that used by [Olsson and Westerborn \(2017\)](#), in which the order of  $g_{t;\theta}$  and  $\mathbf{Q}_\theta$  is swapped.) With this notation, by the filtering recursion (4)–(5),

$$\pi_{t+1;\theta} = \frac{\pi_{t;\theta}\mathbf{L}_{t;\theta}}{\pi_{t;\theta}\mathbf{L}_{t;\theta}\mathbb{1}_X}, \tag{20}$$

with, as previously,  $\phi_{0|-1} := \chi$ . This condensed form of the filtering recursion will be used in Sect. A.3.

In the coming analysis, the following decomposition will be instrumental. For all  $t \in \mathbb{N}$ ,

$$\begin{aligned} \eta_{t;\theta}^N f_t - \eta_{t;\theta} f_t &= \frac{1}{N} \sum_{i=1}^N \left\{ \left( \tau_t^i - \pi_{t;\theta} \mathbf{T}_{t;\theta} h_{t;\theta} \right) \left( f_t(\xi_t^i) - \pi_{t;\theta} f_t \right) \right\} \\ &\quad - \pi_{t;\theta} \left\{ \left( \mathbf{T}_{t;\theta} h_{t;\theta} - \pi_{t;\theta} \mathbf{T}_{t;\theta} h_{t;\theta} \right) \left( f_t - \pi_{t;\theta} f_t \right) \right\} \\ &\quad - \left( \frac{1}{N} \sum_{i=1}^N \tau_t^i - \pi_{t;\theta} \mathbf{T}_{t;\theta} h_{t;\theta} \right) \left( \frac{1}{N} \sum_{i=1}^N f_t(\xi_t^i) - \pi_{t;\theta} f_t \right). \end{aligned} \tag{21}$$

#### A.1 Proof of Theorem 1

We apply the decomposition (21). Note that

$$\begin{aligned} &\frac{1}{N} \sum_{i=1}^N \left\{ \left( \tau_t^i - \pi_{t;\theta} \mathbf{T}_{t;\theta} h_{t;\theta} \right) \left( f_t(\xi_t^i) - \pi_{t;\theta} f_t \right) \right\} \\ &\quad - \pi_{t;\theta} \left\{ \left( \mathbf{T}_{t;\theta} h_{t;\theta} - \pi_{t;\theta} \mathbf{T}_{t;\theta} h_{t;\theta} \right) \left( f_t - \pi_{t;\theta} f_t \right) \right\} \end{aligned}$$



$$\begin{aligned}
 &= \frac{1}{N} \sum_{i=1}^N \omega_t^i \left\{ \tau_t^i \mathfrak{F}_{t;\theta} \left( \xi_t^i \right) + \tilde{\mathfrak{F}}_{t;\theta} \left( \xi_t^i \right) \right\} \\
 &\quad - \pi_{t;\theta} \left\{ g_{t;\theta} \left( \mathbf{T}_{t;\theta} h_{t;\theta} \mathfrak{F}_{t;\theta} + \tilde{\mathfrak{F}}_{t;\theta} \right) \right\},
 \end{aligned} \tag{22}$$

where

$$\begin{aligned}
 \mathfrak{F}_{t;\theta}(x) &:= \frac{f_t(x) - \pi_{t;\theta} f_t}{g_{t;\theta}(x)}, \\
 \tilde{\mathfrak{F}}_{t;\theta}(x) &:= - \frac{\pi_{t;\theta} \mathbf{T}_{t;\theta} h_{t;\theta} \{f_t(x) - \pi_{t;\theta} f_t\}}{g_{t;\theta}(x)},
 \end{aligned} \quad x \in \mathbf{X}. \tag{23}$$

Now, since  $\mathfrak{F}_t g_{t;\theta}$  and  $\tilde{\mathfrak{F}}_t g_{t;\theta}$  both belong to  $F(\mathcal{X})$ , Proposition 1 provides constants  $c_t > 0$  and  $\tilde{c}_t > 0$  such that for all  $\varepsilon > 0$ ,

$$\begin{aligned}
 \mathbb{P} \left( \left| \frac{1}{N} \sum_{i=1}^N \omega_t^i \left\{ \tau_t^i \mathfrak{F}_{t;\theta} \left( \xi_t^i \right) + \tilde{\mathfrak{F}}_{t;\theta} \left( \xi_t^i \right) \right\} - \pi_{t;\theta} \left\{ g_{t;\theta} \left( \mathbf{T}_{t;\theta} h_{t;\theta} \mathfrak{F}_{t;\theta} + \tilde{\mathfrak{F}}_{t;\theta} \right) \right\} \right| \geq \varepsilon \right) \\
 \leq c_t \exp \left( -\tilde{c}_t N \varepsilon^2 \right).
 \end{aligned} \tag{24}$$

To deal with the second part of the decomposition (21), we use the same technique. First, by applying Proposition 1 with  $f \equiv 1/g_{t;\theta}$  and  $\tilde{f} \equiv 0$ , we obtain constants  $a_t > 0$  and  $\tilde{a}_t > 0$  such that for all  $\varepsilon > 0$ ,

$$\mathbb{P} \left( \left| \frac{1}{N} \sum_{i=1}^N \tau_t^i - \pi_{t;\theta} \mathbf{T}_{t;\theta} h_{t;\theta} \right| \geq \varepsilon \right) \leq a_t \exp \left( -\tilde{a}_t N \varepsilon^2 \right). \tag{25}$$

Similarly, using Proposition 1 with  $f \equiv 0$  and  $\tilde{f} \equiv f_t/g_{t;\theta}$  provides constants  $b_t > 0$  and  $\tilde{b}_t > 0$  such that for all  $\varepsilon > 0$ ,

$$\mathbb{P} \left( \left| \frac{1}{N} \sum_{i=1}^N f_t \left( \xi_t^i \right) - \pi_{t;\theta} f_t \right| \geq \varepsilon \right) \leq b_t \exp \left( -\tilde{b}_t N \varepsilon^2 \right). \tag{26}$$

Combining (24), (25), and (26) yields, for all  $\varepsilon > 0$ ,

$$\begin{aligned}
 \mathbb{P} \left( \left| \eta_{t;\theta}^N f_t - \eta_{t;\theta} f_t \right| \geq \varepsilon \right) &\leq a_t \exp \left( -\tilde{a}_t N \varepsilon / 2 \right) + b_t \exp \left( -\tilde{b}_t N \varepsilon / 2 \right) \\
 &\quad + c_t \exp \left\{ -\tilde{c}_t N \left( \varepsilon / 2 \right)^2 \right\},
 \end{aligned}$$

from which the statement of the theorem follows. □

The following result is obtained by inspection of the proof of Olsson and Westerbom (2017, Theorem 1(i)).

**Proposition 1** *Let Assumption 1 hold. Then, for all  $t \in \mathbb{N}$ , all  $\theta \in \Theta$ , all additive state functionals  $h_t \in F(\mathcal{X}^{\otimes(t+1)})$ , all measurable functions  $f_t$  and  $\tilde{f}_t$  such that  $f_t g_{t;\theta} \in$*

$F(\mathcal{X})$  and  $\tilde{f}_t g_{t;\theta} \in F(\mathcal{X})$ , and all  $\tilde{N} \in \mathbb{N}$ , there exist constants  $c_t > 0$  and  $\tilde{c}_t > 0$  (possibly depending on  $\theta, h_t, f_t, \tilde{f}_t$ , and  $\tilde{N}$ ) such that for all  $\varepsilon > 0$ ,

$$\mathbb{P} \left( \left| \frac{1}{N} \sum_{i=1}^N \omega_t^i \left\{ \tau_t^i f_t \left( \xi_t^i \right) + \tilde{f}_t \left( \xi_t^i \right) \right\} - \pi_{t;\theta} \{ g_{t;\theta} (\mathbf{T}_{t;\theta} h_t f_t + \tilde{f}_t) \} \right| \geq \varepsilon \right) \leq c_t \exp(-\tilde{c}_t N \varepsilon^2),$$

where  $\{(\xi_t^i, \omega_t^i, \tau_t^i)\}_{i=1}^N$  are produced using the PaRIS algorithm.

**A.2 Proof of corollary 1**

The  $\mathbb{P}$ -a.s. convergence of  $\eta_{t;\theta}^N f_t$  to  $\eta_{t;\theta} f_t$  is implied straightforwardly by the exponential convergence rate in Theorem 1. Indeed, note that

$$\mathbb{P} \left( \lim_{N \rightarrow \infty} \eta_{t;\theta}^N f_t \neq \eta_{t;\theta} f_t \right) = \lim_{k \rightarrow \infty} \lim_{n \rightarrow \infty} \mathbb{P} \left( \bigcup_{n \leq N} \left\{ \left| \eta_{t;\theta}^N f_t - \eta_{t;\theta} f_t \right| \geq \frac{1}{k} \right\} \right);$$

now, by Theorem 1,

$$\mathbb{P} \left( \bigcup_{n \leq N} \left\{ \left| \eta_{t;\theta}^N f_t - \eta_{t;\theta} f_t \right| \geq \frac{1}{k} \right\} \right) \leq c_t \exp(-\tilde{c}_t n k^{-2}) \sum_{N=0}^{\infty} \exp(-\tilde{c}_t N k^{-2}),$$

where the right-hand side tends to zero when  $n$  tends to infinity. This completes the proof. □

**A.3 Proof of Theorem 2**

By combining (21) and (22),

$$\begin{aligned} & \sqrt{N} \left( \eta_{t;\theta}^N f_t - \eta_{t;\theta} f_t \right) \\ &= \sqrt{N} \left( \frac{1}{N} \sum_{i=1}^N \left\{ \tau_t^i \tilde{\mathfrak{F}}_{t;\theta} \left( \xi_t^i \right) + \tilde{\mathfrak{F}}_{t;\theta} \left( \xi_t^i \right) \right\} - \pi_{t;\theta} (\mathbf{T}_{t;\theta} h_{t;\theta} \tilde{\mathfrak{F}}_{t;\theta} + \tilde{\mathfrak{F}}_{t;\theta}) \right) \\ & \quad - \sqrt{N} \left( \frac{1}{N} \sum_{i=1}^N f_t \left( \xi_t^i \right) - \pi_{t;\theta} f_t \right) \left( \frac{1}{N} \sum_{i=1}^N \tau_t^i - \pi_{t;\theta} \mathbf{T}_{t;\theta} h_{t;\theta} \right), \end{aligned}$$

where in this case

$$\begin{aligned} \tilde{\mathfrak{F}}_{t;\theta}(x) &:= f_t(x) - \pi_{t;\theta} f_t, \\ \tilde{\mathfrak{F}}_{t;\theta}(x) &:= -\pi_{t;\theta} \mathbf{T}_{t;\theta} h_{t;\theta} \{ f_t(x) - \pi_{t;\theta} f_t \}, \end{aligned} \quad x \in \mathcal{X}.$$

are defined in (23). By Proposition 2, since  $\mathfrak{F}_{t;\theta} \in F(X)$  and  $\tilde{\mathfrak{F}}_{t;\theta} \in F(X)$ ,

$$\begin{aligned} & \sqrt{N} \left( \frac{1}{N} \sum_{i=1}^N \left\{ \tau_t^i \mathfrak{F}_{t;\theta}(\xi_t^i) + \tilde{\mathfrak{F}}_{t;\theta}(\xi_t^i) \right\} - \pi_{t;\theta}(\mathbf{T}_{t;\theta} h_{t;\theta} \mathfrak{F}_{t;\theta} + \tilde{\mathfrak{F}}_{t;\theta}) \right) \\ & \xrightarrow{\mathcal{D}} \sigma_{t;\theta} \langle \mathfrak{F}_{t;\theta}, \tilde{\mathfrak{F}}_{t;\theta} \rangle (h_t) Z, \end{aligned}$$

where  $Z$  is standard normally distributed and

$$\sigma_{t;\theta} \langle \mathfrak{F}_{t;\theta}, \tilde{\mathfrak{F}}_{t;\theta} \rangle (h_{t;\theta}) = \sigma_{t;\theta} (h_{t;\theta}),$$

with  $\sigma_{t;\theta} (h_{t;\theta})$  being defined in (14). Now, Proposition 1 and Proposition 2 yield

$$\frac{1}{N} \sum_{i=1}^N \tau_t^i \xrightarrow{\mathbb{P}} \pi_{t;\theta} \mathbf{T}_{t;\theta} h_{t;\theta},$$

and

$$\sqrt{N} \left( \frac{1}{N} \sum_{i=1}^N f_t(\xi_t^i) - \pi_{t;\theta} f_t \right) \xrightarrow{\mathcal{D}} \sigma_{t;\theta}^2 \langle 0, f_t \rangle (h_{t;\theta}) Z$$

(with 0 denoting the zero function), respectively, implying, by Slutsky’s theorem,

$$\sqrt{N} \left( \frac{1}{N} \sum_{i=1}^N f_t(\xi_t^i) - \pi_{t;\theta} f_t \right) \left( \frac{1}{N} \sum_{i=1}^N \tau_t^i - \pi_{t;\theta} \mathbf{T}_{t;\theta} h_{t;\theta} \right) \xrightarrow{\mathbb{P}} 0.$$

Finally, we complete the proof by noting that the term  $\tilde{\sigma}_{t;\theta}^2(f_t)$  in (14) coincides with the asymptotic variance provided by Del Moral et al. (2015, Theorem 3.2).  $\square$

**Proposition 2** *Assumption 1 hold. Then for all  $t \in \mathbb{N}$ , all  $\theta \in \Theta$ , all additive state functionals  $h_t \in F(\mathcal{X}^{\otimes(t+1)})$ , all measurable functions  $f_t \in F(\mathcal{X})$  and  $\tilde{f}_t \in F(\mathcal{X})$ , and all  $\tilde{N} \in \mathbb{N}$ , as  $N \rightarrow \infty$ ,*

$$\sqrt{N} \left( \frac{1}{N} \sum_{i=1}^N \{ \tau_t^i f_t(\xi_t^i) + \tilde{f}_t(\xi_t^i) \} - \pi_{t;\theta}(\mathbf{T}_{t;\theta} h_t f_t + \tilde{f}_t) \right) \xrightarrow{\mathcal{D}} \sigma_{t;\theta} \langle f_t, \tilde{f}_t \rangle (h_t) Z,$$

where  $Z$  is a standard Gaussian random variable and

$$\sigma_{t;\theta}^2 \langle f_t, \tilde{f}_t \rangle (h_t) := \tilde{\sigma}_{t;\theta}^2 \langle f_t, \tilde{f}_t \rangle (h_t) + \sum_{s=0}^{t-1} \sum_{\ell=0}^s \tilde{N}^{\ell-(s+1)} \zeta_{s,\ell,t;\theta} \langle f_t \rangle (h_t), \tag{27}$$

with

$$\sigma_{t;\theta}^2 \langle f_t, \tilde{f}_t \rangle (h_t) := \sum_{s=0}^{t-1} \frac{\pi_{s+1;\theta} \tilde{\mathbf{D}}_{s+1,t;\theta}^2 (h_t f_t + \tilde{f}_t)}{(\pi_{s+1;\theta} \mathbf{L}_{s+1;\theta} \cdots \mathbf{L}_{t-1;\theta} \mathbb{1}_{\mathcal{X}})^2}$$

and

$$\varsigma_{s,\ell,t;\theta} \langle f_t \rangle (h_t) := \frac{\pi_{\ell+1;\theta} \left\{ \overleftarrow{\mathbf{Q}}_{\phi_{\ell;\theta}} (\mathbf{T}_{\ell;\theta} h_{\ell} + \tilde{h}_{\ell} - \mathbf{T}_{\ell+1;\theta} h_{\ell+1})^2 \mathbf{L}_{\ell+1;\theta} \cdots \mathbf{L}_{s;\theta} (\mathbf{L}_{s+1;\theta} \cdots \mathbf{L}_{t-1;\theta} f_t)^2 \right\}}{(\pi_{\ell+1;\theta} \mathbf{L}_{\ell+1;\theta} \cdots \mathbf{L}_{s;\theta} \mathbb{1}_{\mathcal{X}}) (\pi_{s+1;\theta} \mathbf{L}_{s+1;\theta} \cdots \mathbf{L}_{t-1;\theta} \mathbb{1}_{\mathcal{X}})^2}.$$

**Proof of Proposition 2** Assume first that  $\pi_{t;\theta} (\mathbf{T}_{t;\theta} h_t f_t + \tilde{f}_t) = 0$ . Then, by Lemma 1 and Slutsky’s theorem, as  $\Omega_t/N \xrightarrow{\mathbb{P}} \pi_{t;\theta} g_{t;\theta}$  by Proposition 1,

$$\begin{aligned} & \sqrt{N} \left( \frac{1}{N} \sum_{i=1}^N \{ \tau_t^i f_t(\xi_t^i) + \tilde{f}_t(\xi_t^i) \} \right) \\ &= \sqrt{N} \left( \sum_{i=1}^N \frac{\omega_t^i}{\Omega_t} \{ \tau_t^i \mathfrak{F}_{t;\theta}(\xi_t^i) + \tilde{\mathfrak{F}}_{t;\theta}(\xi_t^i) \} \right) \frac{1}{N} \Omega_t \xrightarrow{\mathcal{D}} \pi_{t;\theta} g_{t;\theta} \times \Gamma_{t;\theta} \langle \mathfrak{F}_{t;\theta}, \tilde{\mathfrak{F}}_{t;\theta} \rangle (h_t) Z, \end{aligned}$$

where again  $Z$  has standard Gaussian distribution,  $\Gamma_{t;\theta} \langle \mathfrak{F}_{t;\theta}, \tilde{\mathfrak{F}}_{t;\theta} \rangle (h_t)$  is given in Lemma 1, and we have set  $\mathfrak{F}_{t;\theta} := f_t/g_{t;\theta}$  and  $\tilde{\mathfrak{F}}_{t;\theta} := \tilde{f}_t/g_{t;\theta}$  and used, first, that  $g_{t;\theta} \mathfrak{F}_{t;\theta} = f_t \in F(\mathcal{X})$  and  $g_{t;\theta} \tilde{\mathfrak{F}}_{t;\theta} = \tilde{f}_t \in F(\mathcal{X})$  and, second, that

$$\phi_{t;\theta} (\mathbf{T}_{t;\theta} h_t \mathfrak{F}_{t;\theta} + \tilde{\mathfrak{F}}_{t;\theta}) = \frac{\pi_{t;\theta} \{ g_{t;\theta} (\mathbf{T}_{t;\theta} h_t \mathfrak{F}_{t;\theta} + \tilde{\mathfrak{F}}_{t;\theta}) \}}{\pi_{t;\theta} g_{t;\theta}} = \frac{\pi_{t;\theta} (\mathbf{T}_{t;\theta} h_t f_t + \tilde{f}_t)}{\pi_{t;\theta} g_{t;\theta}} = 0.$$

Now, by iterating (20) we conclude that for all  $(s, t) \in \mathbb{N}^2$ ,

$$\begin{aligned} \pi_{t;\theta} &= \frac{\pi_{t-1;\theta} \mathbf{L}_{t-1;\theta}}{\pi_{t-1;\theta} \mathbf{L}_{t-1;\theta} \mathbb{1}_{\mathcal{X}}} = \frac{\pi_{t-2;\theta} \mathbf{L}_{t-2;\theta} \mathbf{L}_{t-1;\theta}}{(\pi_{t-2;\theta} \mathbf{L}_{t-2;\theta} \mathbb{1}_{\mathcal{X}}) (\pi_{t-1;\theta} \mathbf{L}_{t-1;\theta} \mathbb{1}_{\mathcal{X}})} \\ &= \frac{\pi_{t-2;\theta} \mathbf{L}_{t-2;\theta} \mathbf{L}_{t-1;\theta}}{\pi_{t-2;\theta} \mathbf{L}_{t-2;\theta} \mathbf{L}_{t-1;\theta} \mathbb{1}_{\mathcal{X}}} \\ &= \frac{\pi_{s+1;\theta} \mathbf{L}_{s+1;\theta} \cdots \mathbf{L}_{t-1;\theta}}{\pi_{s+1;\theta} \mathbf{L}_{s+1;\theta} \cdots \mathbf{L}_{t-1;\theta} \mathbb{1}_{\mathcal{X}}} \end{aligned}$$

and, consequently,

$$\pi_{t;\theta} g_{t;\theta} = \frac{\pi_{s+1;\theta} \mathbf{L}_{s+1;\theta} \cdots \mathbf{L}_{t-1;\theta} g_{t;\theta}}{\pi_{s+1;\theta} \mathbf{L}_{s+1;\theta} \cdots \mathbf{L}_{t-1;\theta} \mathbb{1}_{\mathcal{X}}}. \tag{28}$$

Finally, by (28) it holds that

$$\pi_{t;\theta} g_{t;\theta} \times \Gamma_{t;\theta} \langle \mathfrak{F}_{t;\theta}, \tilde{\mathfrak{F}}_{t;\theta} \rangle (h_t) = \sigma_{t;\theta} \langle f_t, \tilde{f}_t \rangle (h_t),$$

where  $\Gamma_{t;\theta}^2 \langle f_t, \tilde{f}_t \rangle (h_t)$  is defined in (29) and  $\sigma_{t;\theta}^2 \langle f_t, \tilde{f}_t \rangle (h_t)$  is defined in (27). Finally, in the general case, the previous holds true when  $\tilde{f}_t$  is replaced by  $\tilde{f}_t - \pi_{t;\theta}(\mathbf{T}_{t;\theta} h_t f_t + \tilde{f}_t)$ , which completes the proof.

The following lemma is obtained by inspection of the proof of Olsson and Westerbom (2017, Theorem 3).

**Lemma 1** *Assumption 1 hold. Then for all  $t \in \mathbb{N}$ , all  $\theta \in \Theta$ , all additive state functionals  $h_t \in F(\mathcal{X}^{\otimes(t+1)})$ , all measurable functions  $f_t$  and  $\tilde{f}_t$  such that  $f_t g_{t;\theta} \in F(\mathcal{X})$  and  $\tilde{f}_t g_{t;\theta} \in F(\mathcal{X})$ , and all  $\tilde{N} \in \mathbb{N}$ , as  $N \rightarrow \infty$ ,*

$$\sqrt{N} \left( \sum_{i=1}^N \frac{\omega_i}{\Omega_t} \{ \tau_t^i f_t(\xi_t^i) + \tilde{f}_t(\xi_t^i) \} - \phi_{t;\theta}(\mathbf{T}_{t;\theta} h_t f_t + \tilde{f}_t) \right) \xrightarrow{\mathcal{D}} \Gamma_{t;\theta} \langle f_t, \tilde{f}_t \rangle (h_t) Z,$$

where  $Z$  is a standard normal distribution and

$$\Gamma_{t;\theta}^2 \langle f_t, \tilde{f}_t \rangle (h_t) := \tilde{\Gamma}_{t;\theta}^2 \langle f_t, \tilde{f}_t \rangle (h_t) + \sum_{s=0}^{t-1} \sum_{\ell=0}^s \tilde{N}^{\ell-(s+1)} \gamma_{s,\ell,t;\theta} \langle f_t \rangle (h_t), \tag{29}$$

with

$$\tilde{\Gamma}_{t;\theta}^2 \langle f_t, \tilde{f}_t \rangle (h_t) := \sum_{s=0}^{t-1} \frac{\pi_{s+1;\theta} \tilde{\mathbf{D}}_{s+1,t;\theta}^2 \{ g_{t;\theta}(h_t f_t + \tilde{f}_t) \}}{(\pi_{s+1;\theta} \mathbf{L}_{s+1;\theta} \cdots \mathbf{L}_{t-1;\theta} g_{t;\theta})^2}$$

and

$$\gamma_{s,\ell,t;\theta} \langle f_t \rangle (h_t) := \frac{\pi_{\ell+1;\theta} \left\{ \tilde{\mathbf{Q}}_{\phi_{\ell;\theta}}(\mathbf{T}_{\ell;\theta} h_\ell + \tilde{h}_\ell - \mathbf{T}_{\ell+1;\theta} h_{\ell+1})^2 \mathbf{L}_{\ell+1;\theta} \cdots \mathbf{L}_{s;\theta} (\mathbf{L}_{s+1;\theta} \cdots \mathbf{L}_{t-1;\theta} g_{t;\theta} f_t)^2 \right\}}{(\pi_{\ell+1;\theta} \mathbf{L}_{\ell+1;\theta} \cdots \mathbf{L}_{s;\theta} \mathbb{1}_X) (\pi_{s+1;\theta} \mathbf{L}_{s+1;\theta} \cdots \mathbf{L}_{t-1;\theta} g_{t;\theta})^2}.$$

### A.4 Proof of Theorem 3

As noted above, the first term of the asymptotic variance  $\sigma_t^2(f_t)$  coincides with the asymptotic variance  $\tilde{\sigma}_{t;\theta}^2(f_t)$  obtained by Del Moral et al. (2015, Theorem 3.2). The same work provides a constant  $c \in \mathbb{R}_+$  such that  $\tilde{\sigma}_{t;\theta}^2(f_t) \leq c \|f_t\|_\infty$ , and we may hence focus on bounding second term of the asymptotic variance.

For this purpose, note that for all  $s \leq t - 1$  and  $x_{s+1} \in X$ ,

$$\begin{aligned} & \mathbf{L}_{s+1;\theta} \cdots \mathbf{L}_{t-1;\theta} (f_t - \pi_{t;\theta} f_t)(x_{s+1}) \\ &= g_{s+1;\theta}(x_{s+1}) \mathbf{Q}_\theta \mathbf{L}_{s+2;\theta} \cdots \mathbf{L}_{t-2;\theta} \{ g_{t-1;\theta} (\mathbf{Q}_\theta f_t - \phi_{t-1;\theta} \mathbf{Q}_\theta f_t) \} (x_{s+1}). \end{aligned} \tag{30}$$

By applying the forgetting of the filter, or, more particularly, Douc et al. (2011, Lemma 10), to (30) we obtain

$$\| \mathbf{L}_{s+1;\theta} \cdots \mathbf{L}_{t-1;\theta} (f_t - \pi_{t;\theta} f_t) \|_\infty \leq 2\bar{\delta} \| \mathbf{Q}_\theta \mathbf{L}_{s+2;\theta} \cdots \mathbf{L}_{t-1;\theta} \mathbb{1}_X \|_\infty \| f_t \|_\infty^2 e^{t-s-2}.$$

Note that in the previous bound, the exponential contraction follows from the fact that the objective function  $f_t$  is centred around its predicted mean. The latter is a consequence of the fact that the tangent filter is, as a covariance, centred itself (recall the identities (10) and (11) and the decomposition (21)). In addition, from the proof of Olsson and Westerborn (2017, Theorem 8) we extract, using Assumption 3,

$$\|\mathbf{T}_{\ell;\theta} h_{\ell;\theta} + \tilde{h}_{\ell;\theta} - \mathbf{T}_{\ell+1;\theta} h_{\ell+1;\theta}\|_{\infty} \leq |\tilde{h}|_{\infty} \left( 4\bar{\delta} \frac{\varrho^2}{1-\varrho} + 1 \right),$$

and under Assumption 2, for all  $x \in X$ ,

$$\rho^{-1} \bar{\delta} \mu_{\mathbf{L}_{s+2;\theta}} \cdots \mathbf{L}_{t-1;\theta} \mathbb{1}_X \leq \mathbf{L}_{s+1;\theta} \cdots \mathbf{L}_{t-1;\theta} \mathbb{1}_X(x) \leq \rho \bar{\delta} \mu_{\mathbf{L}_{s+2;\theta}} \cdots \mathbf{L}_{t-1;\theta} \mathbb{1}_X.$$

Combining the previous bounds gives

$$\begin{aligned} s_{s,\ell,t;\theta}(f_t) &= \frac{\pi_{\ell+1;\theta} \left\{ \bar{\mathbf{Q}}_{\phi_{\ell;\theta}} (\mathbf{T}_{\ell;\theta} h_{\ell;\theta} + \tilde{h}_{\ell;\theta} - \mathbf{T}_{\ell+1;\theta} h_{\ell+1;\theta})^2 \mathbf{L}_{\ell+1;\theta} \cdots \mathbf{L}_{s;\theta} (\mathbf{L}_{s+1;\theta} \cdots \mathbf{L}_{t-1;\theta} (f_t - \pi_{t;\theta} f_t))^2 \right\}}{(\pi_{\ell+1;\theta} \mathbf{L}_{\ell+1;\theta} \cdots \mathbf{L}_{s;\theta} \mathbb{1}_X) (\pi_{s+1;\theta} \mathbf{L}_{s+1;\theta} \cdots \mathbf{L}_{t-1;\theta} \mathbb{1}_X)^2} \\ &\leq \tilde{c} \|f_t\|_{\infty}^2 \varrho^{2(t-s-2)}, \end{aligned}$$

where

$$\tilde{c} := 4|\tilde{h}|_{\infty}^2 \left( 4\bar{\delta} \frac{\varrho^2}{1-\varrho} + 1 \right)^2 \left( \frac{\bar{\delta}}{\underline{\delta}(1-\varrho)} \right)^3.$$

Finally, summing up yields

$$\begin{aligned} \sum_{s=0}^{t-1} \sum_{\ell=0}^s s_{s,\ell,t;\theta}(f_t) &\leq \tilde{c} \|f_t\|_{\infty}^2 \sum_{s=0}^{t-1} \sum_{\ell=0}^s \tilde{N}^{\ell-s-1} \varrho^{2(t-s-2)} \\ &= \tilde{c} \|f_t\|_{\infty}^2 \frac{1}{\tilde{N}-1} \sum_{s=0}^{t-1} \varrho^{2(t-s-2)} (1 - \tilde{N}^{-s-1}) \\ &\leq \tilde{c} \|f_t\|_{\infty}^2 \frac{1}{\tilde{N}-1} \sum_{s=0}^{t-1} \varrho^{2(t-s-2)} \leq \tilde{c} \|f_t\|_{\infty}^2 \frac{1}{(\tilde{N}-1)(1-\varrho^2)}, \end{aligned}$$

which completes the proof. □

### B Kernels

Given two measurable spaces  $(X, \mathcal{X})$  and  $(Y, \mathcal{Y})$ , an unnormalised transition kernel  $\mathbf{K}$  between these spaces induces two integral operators, one acting on functions and the other on measures. Specifically, we define the function

$$\mathbf{K}h : X \ni x \mapsto \int h(x, y) \mathbf{K}(x, dy) \quad (h \in F(\mathcal{X} \otimes \mathcal{Y}))$$

and the measure

$$\nu \mathbf{K} : \mathcal{Y} \ni A \mapsto \int \mathbf{K}(x, A) \nu(dx) \quad (\nu \in \mathcal{M}(\mathcal{X}))$$

whenever these quantities are well defined. Moreover, let  $\mathbf{L}$  be another unnormalised transition kernel from  $(Y, \mathcal{Y})$  to the measurable space  $(Z, \mathcal{Z})$ ; then two different *products* of  $\mathbf{K}$  and  $\mathbf{L}$  can be defined, namely

$$\mathbf{KL} : X \times \mathcal{Z} \ni (x, A) \mapsto \int \mathbf{K}(x, dy) \mathbf{L}(y, A)$$

and

$$\mathbf{K} \otimes \mathbf{L} : X \times (\mathcal{Y} \otimes \mathcal{Z}) \ni (x, A) \mapsto \int \mathbb{1}_A(y, z) \mathbf{K}(x, dy) \mathbf{L}(y, dz)$$

whenever these are well defined. These products form new transition kernels from  $(X, \mathcal{X})$  to  $(Z, \mathcal{Z})$  and from  $(X, \mathcal{X})$  to  $(Y \times Z, \mathcal{Y} \otimes \mathcal{Z})$ , respectively. Also the  $\otimes$ -product of a kernel  $\mathbf{K}$  and a measure  $\nu \in \mathcal{M}(\mathcal{X})$  is defined as the new measure

$$\nu \otimes \mathbf{K} : \mathcal{X} \otimes \mathcal{Y} \ni A \mapsto \int \mathbb{1}_A(x, y) \mathbf{K}(x, dy) \nu(dx).$$

Finally, for any kernel  $\mathbf{K}$  and any bounded measurable function  $h$  we write  $\mathbf{K}^2 h := (\mathbf{K}h)^2$  and  $\mathbf{K}h^2 := \mathbf{K}(h^2)$ . Similar notation will be used for measures.

## References

- Anderson, B. D. O., Moore, J. B. (1979). *Optimal filtering*. New Jersey: Prentice-Hall.
- Cappé, O. (2001). Ten years of HMMs (online bibliography 1989–2000). <http://perso.telecom-paristech.fr/~cappe/docs/hmmbib.html>. Accessed Mar 2013.
- Cappé, O. (2011). Online EM algorithm for hidden Markov models. *Journal of Computational and Graphical Statistics*, 20(3), 728–749.
- Cappé, O., Moulines, E., Rydén, T. (2005). *Inference in hidden Markov models*. New York: Springer.
- Crisan, D., Heine, K. (2008). Stability of the discrete time filter in terms of the tails of noise distributions. *Journal of the London Mathematical Society*, 78(2), 441–458.
- Del Moral, P., Guionnet, A. (2001). On the stability of interacting processes with applications to filtering and genetic algorithms. *Annales de l'Institut Henri Poincaré*, 37(2), 155–194.
- Del Moral, P., Doucet, A., Singh, S. (2010). A backward particle interpretation of Feynman–Kac formulae. *ESAIM: Mathematical Modelling and Numerical Analysis*, 44(5), 947–975.
- Del Moral, P., Doucet, A., Singh, S. S. (2015). Uniform stability of a particle approximation of the optimal filter derivative. *SIAM Journal on Control and Optimization*, 53(3), 1278–1304.
- Delyon, B. (1996). General results on the convergence of stochastic algorithms. *IEEE Transactions on Automatic Control*, 41(9), 1245–1255.
- Dempster, A. P., Laird, N. M., Rubin, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society*, 39(1), 1–38 (with discussion).
- Douc, R., Matias, C. (2001). Asymptotics of the maximum likelihood estimator for general hidden Markov models. *Bernoulli*, 7(3), 381–420.
- Douc, R., Garivier, A., Moulines, E., Olsson, J. (2011). Sequential Monte Carlo smoothing for general state space hidden Markov models. *Annals of Applied Probability*, 21(6), 2109–2145.
- Douc, R., Moulines, E., Olsson, J. (2014). Long-term stability of sequential Monte Carlo methods under verifiable conditions. *Annals of Applied Probability*, 24(5), 1767–1802.

- Doucet, A., Tadić, V. B. (2003). Parameter estimation in general state-space models using particle methods. *Annals of the Institute of Statistical Mathematics*, 55(2), 409–422.
- Doucet, A., Godsill, S., Andrieu, C. (2000). On sequential Monte-Carlo sampling methods for Bayesian filtering. *Statistics and Computing*, 10(3), 197–208.
- Doucet, A., De Freitas, N., Gordon, N. (Eds.). (2001). *Sequential Monte Carlo methods in practice*. New York: Springer.
- Fearnhead, P., Wyncoll, D., Tawn, J. (2010). A sequential smoothing algorithm with linear computational cost. *Biometrika*, 97(2), 447–464.
- Gordon, N., Salmond, D., Smith, A. F. (1993). Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEE Proceedings F-Radar and Signal Processing*, 140(2), 107–113.
- Hull, J., White, A. (1987). The pricing of options on assets with stochastic volatilities. *The Journal of Finance*, 42(2), 281–300.
- Jacob, P. E., Murray, L. M., Rubenthaler, S. (2013). Path storage in the particle filter. *Statistics and Computing*, 25(2), 487–496.
- Jasra, A. (2015). On the behaviour of the backward interpretation of Feynman–Kac formulae under verifiable conditions. *Journal of Applied Probability*, 52(2), 339–359.
- Julier, S. J., Uhlmann, J. K. (1997). A new extension of the Kalman filter to nonlinear systems. In *AeroSense: The 11th international symposium on aerospace/defense sensing. Simulation and controls*.
- Kantas, N., Doucet, A., Singh, S. S., Maciejowski, J., Chopin, N. (2015). On particle methods for parameter estimation in state-space models. *Statistical Science*, 30(3), 328–351.
- Kitagawa, G. (1996). Monte Carlo filter and smoother for non-Gaussian nonlinear state space models. *Journal of Computational and Graphical Statistics*, 5(1), 1–25.
- Kitagawa, G., Sato, S. (2001). Monte Carlo smoothing and self-organising state-space model. In *Sequential Monte Carlo methods in practice* (pp. 177–195). New York: Springer.
- Le Corff, S., Fort, G., Moulines, E. (2011). Online expectation maximization algorithm to solve the SLAM problem. In *2011 IEEE statistical signal processing workshop (SSP)* (pp. 225–228).
- Le Gland, F., Mevel, L. (1996). Geometric ergodicity in hidden markov models. In *Research report, RR-2991, INRIA*.
- Le Gland, F., Mevel, L. (1997) Recursive estimation in hidden Markov models. In *Proceedings of the 36th IEEE conference on decision and control* (pp. 3468–3473).
- Martinez-Cantin, R., de Freitas, N., Castellanos, J. A. (2007). Analysis of particle methods for simultaneous robot localization and mapping and a new algorithm: Marginal-slam. In *Proceedings 2007 IEEE international conference on robotics and automation* (pp. 2415–2420).
- Montemerlo, M., Thrun, S., Koller, D., Wegbreit, B. (2002). FastSLAM: A factored solution to the simultaneous localization and mapping problem. In *Proceedings of the AAAI national conference on artificial intelligence*. Edmonton: AAAI.
- Montemerlo, M., Thrun, S., Koller, D., Wegbreit, B. (2003). An improved particle filtering algorithm for simultaneous localization and mapping that provably converges. In *Proceedings of the sixteenth international joint conference on artificial intelligence (IJCAI)*. Acapulco: IJCAI.
- Nguyen, T. N. M., Le Corff, S., Moulines, E. (2017). On the two-filter approximations of marginal smoothing distributions in general state-space models. *Advances in Applied Probability*, 50(1), 154–177.
- Olsson, J., Cappé, O., Douc, R., Moulines, E. (2008). Sequential Monte Carlo smoothing with application to parameter estimation in non-linear state space models. *Bernoulli*, 14(1), 155–179.
- Olsson, J., Westerborn, J. (2016). Efficient parameter inference in general hidden markov models using the filter derivatives. In *2016 IEEE international conference on acoustics, speech and signal processing (ICASSP)* (pp. 3984–3988).
- Olsson, J., Westerborn, J. (2017). Efficient particle-based online smoothing in general hidden Markov models: The PaRIS algorithm. *Bernoulli*, 23(3), 1951–1996.
- Poyiadjis, G., Doucet, A., Singh, S. (2011). Particle approximations of the score and observed information matrix in state space models with application to parameter estimation. *Biometrika*, 98(1), 65–80.
- Poyiadjis, G., Doucet, A., Singh, S. S. (2005). Particle methods for optimal filter derivative: application to parameter estimation. In *Proceedings IEEE international conference on acoustics, speech, and signal processing* (pp. 925–928).
- Tadić, V. B. (2010). Analyticity, convergence, and convergence rate of recursive maximum-likelihood estimation in hidden markov models. *IEEE Transactions on Information Theory*, 56(12), 6406–6432.
- Tadić, V. B., Doucet, A. (2017). Asymptotic bias of stochastic gradient search. *Annals of Applied Probability*, 27(6), 3255–3304.