



**HAL**  
open science

# Coupling stochastic EM and Approximate Bayesian computation for parameter inference in state-space models

Umberto Picchini, Adeline Samson

► **To cite this version:**

Umberto Picchini, Adeline Samson. Coupling stochastic EM and Approximate Bayesian computation for parameter inference in state-space models. *Computational Statistics*, 2018, 33 (1), pp.179-212. 10.1007/s00180-017-0770-y . hal-01623737v2

**HAL Id: hal-01623737**

**<https://hal.science/hal-01623737v2>**

Submitted on 8 Dec 2017

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

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

# Coupling stochastic EM and Approximate Bayesian Computation for parameter inference in state-space models

Umberto Picchini<sup>a,\*</sup>, Adeline Samson<sup>b,\*</sup>

<sup>a</sup>Centre for Mathematical Sciences, Sölvegatan 18, SE-22100 Lund, Sweden

Email: [umberto@maths.lth.se](mailto:umberto@maths.lth.se)

<sup>b</sup>LJK, Universite Grenoble Alpes, F-38000 Grenoble, France;

CNRS, LJK, F-38000 Grenoble, France

Email: [adeline.leclercq-samson@imag.fr](mailto:adeline.leclercq-samson@imag.fr)

## Abstract

We study the class of state-space models and perform maximum likelihood estimation for the model parameters. We consider a stochastic approximation expectation-maximization (SAEM) algorithm to maximize the likelihood function with the novelty of using approximate Bayesian computation (ABC) within SAEM. The task is to provide each iteration of SAEM with a filtered state of the system, and this is achieved using an ABC sampler for the hidden state, based on sequential Monte Carlo (SMC) methodology. It is shown that the resulting SAEM-ABC algorithm can be calibrated to return accurate inference, and in some situations it can outperform a version of SAEM incorporating the bootstrap filter. Two simulation studies are presented, first a nonlinear Gaussian state-space model then a state-space model having dynamics expressed by a stochastic differential equation. Comparisons with iterated filtering for maximum likelihood inference, and Gibbs sampling and particle marginal methods for Bayesian inference are presented.

**Keywords:** hidden Markov model; maximum likelihood; particle filter; SAEM; sequential Monte Carlo; stochastic differential equation.

## 1 Introduction

State-space models [Cappé et al., 2005] are widely applied in many fields, such as biology, chemistry, ecology, etc. Let us now introduce some notation. Consider an observable, discrete-time stochastic process  $\{\mathbf{Y}_t\}_{t \geq t_0}$ ,  $\mathbf{Y}_t \in \mathbf{Y} \subseteq \mathbb{R}^{d_y}$  and a latent and unobserved continuous-time stochastic process  $\{\mathbf{X}_t\}_{t \geq t_0}$ ,  $\mathbf{X}_t \in \mathbf{X} \subseteq \mathbb{R}^{d_x}$ . Process  $\mathbf{X}_t \sim p(\mathbf{x}_t | \mathbf{x}_s, \boldsymbol{\theta}_x)$  is assumed Markov with transition densities  $p(\cdot)$ ,  $s < t$ . Processes  $\{\mathbf{X}_t\}$  and  $\{\mathbf{Y}_t\}$  depend on their own (unknown) vector-parameters  $\boldsymbol{\theta}_x$  and  $\boldsymbol{\theta}_y$ , respectively. We consider  $\{\mathbf{Y}_t\}$  as a measurement-error-corrupted version of  $\{\mathbf{X}_t\}$  and assume that observations for  $\{\mathbf{Y}_t\}$  are conditionally independent given  $\{\mathbf{X}_t\}$ . The state-space model can be summarised as

$$\begin{cases} \mathbf{Y}_t \sim f(\mathbf{y}_t | \mathbf{X}_t, \boldsymbol{\theta}_y), \\ \mathbf{X}_t \sim p(\mathbf{x}_t | \mathbf{x}_s, \boldsymbol{\theta}_x), \quad t_0 \leq s < t, \quad \mathbf{X}_0 \sim p(\mathbf{x}_0) \end{cases} \quad (1)$$

where  $\mathbf{X}_0 \equiv \mathbf{X}_{t_0}$ . We assume  $f(\cdot)$  a known density (or probability mass) function set by the modeller. Regarding the transition density  $p(\mathbf{x}_t | \mathbf{x}_s, \cdot)$ , this is typically unknown except for very simple toy models.

Goal of our work is to estimate the parameters  $(\boldsymbol{\theta}_x, \boldsymbol{\theta}_y)$  by maximum likelihood, using observations  $\mathbf{Y}_{1:n} = (\mathbf{Y}_1, \dots, \mathbf{Y}_n)$  collected at discrete times  $\{t_1, \dots, t_n\}$ . Here  $\mathbf{Y}_j \equiv \mathbf{Y}_{t_j}$  and we

use  $\mathbf{z}_{1:n}$  to denote a generic sequence  $(\mathbf{z}_1, \dots, \mathbf{z}_n)$ . For ease of notation we refer to the vector  $\boldsymbol{\theta} := (\boldsymbol{\theta}_x, \boldsymbol{\theta}_y)$  as the object of our inference.

Parameters inference for state-space models has been widely developed, and sequential Monte Carlo (SMC) methods are now considered the state-of-art when dealing with nonlinear/non-Gaussian state space models (see [Kantas et al., 2015](#) for a review). Methodological advancements have especially considered Bayesian approaches. In Bayesian inference the goal is to derive analytically the posterior distribution  $\pi(\boldsymbol{\theta}|\mathbf{Y}_{1:n})$  or, most frequently, implement an algorithm for sampling draws from the posterior. Sampling procedures are often carried out using Markov chain Monte Carlo (MCMC) or SMC embedded in MCMC procedures, see [Andrieu and Roberts \[2009\]](#) and [Andrieu et al. \[2010\]](#).

In this work we instead aim at maximum likelihood estimation of  $\boldsymbol{\theta}$ . Several methods for maximum likelihood inference in state-space models have been proposed in the literature, including the well-known EM algorithm [[Dempster et al., 1977](#)]. The EM algorithm computes the conditional expectation of the complete-likelihood for the pair  $(\{\mathbf{Y}_t\}, \{\mathbf{X}_t\})$  and then produces a (local) maximizer for the likelihood function based on the actual observations  $\mathbf{Y}_{1:n}$ . One of the difficulties is how to compute the conditional expectation of the state  $\{\mathbf{X}_t\}$  given the observations  $\mathbf{Y}_{1:n}$ . This conditional expectation can be computed exactly with the Kalman filter when the state-space is linear and Gaussian [[Cappé et al., 2005](#)], otherwise it has to be approximated. In this work we focus on a stochastic approximation of  $\{\mathbf{X}_t\}$ . Therefore, we resort to a stochastic version of the EM algorithm, namely the Stochastic Approximation EM (SAEM) [[Delyon et al., 1999](#)]. The problem is to generate, conditionally on the current value of  $\boldsymbol{\theta}$  during the EM maximization, an appropriate “proposal” for the state  $\{\mathbf{X}_t\}$ , and we use SMC to obtain such proposal. SMC algorithms [[Doucet et al., 2001](#)] have already been coupled to stochastic EM algorithms (see e.g. [Huys et al. \[2006\]](#), [Huys and Paninski \[2009\]](#), [Lindsten \[2013\]](#), [Ditlevsen and Samson \[2014\]](#) and references therein). The simplest and most popular SMC algorithm, the bootstrap filter [[Gordon et al., 1993](#)], is easy to implement and very general, explicit knowledge of the density  $f(\mathbf{y}_t|\mathbf{X}_t, \cdot)$  being the only requirement. Therefore the bootstrap filter is often a go-to option for practitioners. Alternatively, in order to select a path  $\{\mathbf{X}_t\}$  to feed SAEM with, in this paper we follow an approach based on approximate Bayesian computation (ABC) and specifically we use the ABC-SMC method for state-estimation proposed in [Jasra et al. \[2012\]](#). We do not merely consider the algorithm by Jasra et al. within SAEM, but show in detail and discuss how SAEM-ABC-SMC (shortly SAEM-ABC) can in some cases outperform SAEM coupled with the bootstrap filter.

We illustrate our SAEM-ABC approach for (approximate) maximum likelihood estimation of  $\boldsymbol{\theta}$  using two case studies, a nonlinear Gaussian state-space model and a more complex state-space model based on stochastic differential equations. We also compare our method with the iterated filtering for maximum likelihood estimation [[Ionides et al., 2015](#)], Gibbs sampling and particle marginal methods for Bayesian inference [[Andrieu and Roberts, 2009](#)] and will also use a special SMC proposal function for the specific case of stochastic differential equations [[Golightly and Wilkinson, 2011](#)].

The paper is structured as follows: in section 2 we introduce the standard SAEM algorithm and basic notions of ABC. In section 3 we propose a new method for maximum likelihood estimation by integrating an ABC-SMC algorithm within SAEM. Section 4 shows simulation results and section 5 summarize conclusions. An appendix includes technical details pertaining the simulation studies. Software code can be found online either at <https://github.com/umbertopicchini/SAEM-ABC> or as supplementary material in the version of this paper published on *Computational Statistics* (doi:10.1007/s00180-017-0770-y)<sup>1</sup>.

<sup>1</sup>The version of the code available on *Computational Statistics* is the one submitted at the time of the article acceptance. The one on GitHub might contain more recent amendments (if any).

## 2 The complete likelihood and stochastic approximation EM

Recall that  $\mathbf{Y}_{1:n} = (\mathbf{Y}_1, \dots, \mathbf{Y}_n)$  denotes the available data collected at times  $(t_1, \dots, t_n)$  and denote with  $\mathbf{X}_{1:n} = (\mathbf{X}_1, \dots, \mathbf{X}_n)$  the corresponding unobserved states. We additionally set  $\mathbf{X}_{0:n} = (\mathbf{X}_0, \mathbf{X}_{1:n})$  for the vector including an initial (fixed or random) state  $\mathbf{X}_0$ , that is  $\mathbf{X}_1$  is generated as  $\mathbf{X}_1 \sim p(\mathbf{x}_1|\mathbf{x}_0)$ . When the transition densities  $p(\mathbf{x}_j|\mathbf{x}_{j-1})$  are available in closed form ( $j = 1, \dots, n$ ), the likelihood function for  $\boldsymbol{\theta}$  can be written as (here we have assumed a random initial state with density  $p(\mathbf{X}_0)$ )

$$\begin{aligned} p(\mathbf{Y}_{1:n}; \boldsymbol{\theta}) &= \int p_{\mathbf{Y}, \mathbf{X}}(\mathbf{Y}_{1:n}, \mathbf{X}_{0:n}; \boldsymbol{\theta}) d\mathbf{X}_{0:n} = \int p_{\mathbf{Y}|\mathbf{X}}(\mathbf{Y}_{1:n}|\mathbf{X}_{0:n}; \boldsymbol{\theta}) p_{\mathbf{X}}(\mathbf{X}_{0:n}; \boldsymbol{\theta}) d\mathbf{X}_{0:n} \\ &= \int p(\mathbf{X}_0) \left\{ \prod_{j=1}^n f(\mathbf{Y}_j|\mathbf{X}_j; \boldsymbol{\theta}) p(\mathbf{X}_j|\mathbf{X}_{j-1}; \boldsymbol{\theta}) \right\} d\mathbf{X}_0 \cdots d\mathbf{X}_n \end{aligned} \quad (2)$$

where  $p_{\mathbf{Y}, \mathbf{X}}$  is the ‘‘complete data likelihood’’,  $p_{\mathbf{Y}|\mathbf{X}}$  is the conditional law of  $\mathbf{Y}$  given  $\mathbf{X}$ ,  $p_{\mathbf{X}}$  is the law of  $\mathbf{X}$ ,  $f(\mathbf{Y}_j|\mathbf{X}_j; \cdot)$  is the conditional density of  $\mathbf{Y}_j$  as in (1) and  $p_{\mathbf{X}}(\mathbf{X}_{0:n}; \cdot)$  the joint density of  $\mathbf{X}_{0:n}$ . The last equality in (2) exploits the notion of conditional independence of observations given latent states and the Markovian property of  $\{\mathbf{X}_t\}$ . In general the likelihood (2) is not explicitly known either because the integral is multidimensional and because expressions for transition densities are typically not available except for trivial toy models.

In addition, when an exact simulator for the solution of the dynamical process associated with the Markov process  $\{\mathbf{X}_t\}$  is unavailable, hence it is not possible to sample from  $p(\mathbf{X}_j|\mathbf{X}_{j-1}; \boldsymbol{\theta})$ , numerical discretization methods are required. Without loss of generality, say that we have equispaced sampling times such that  $t_j = t_{j-1} + \Delta$ , with  $\Delta > 0$ . Now introduce a discretization for the interval  $[t_1, t_n]$  given by  $\{\tau_1, \tau_h, \dots, \tau_{Rh}, \dots, \tau_{nRh}\}$  where  $h = \Delta/R$  and  $R \geq 1$ . We take  $\tau_1 = t_1$ ,  $\tau_{nRh} = t_n$  and therefore  $\tau_i \in \{t_1, \dots, t_n\}$  for  $i = 1, Rh, 2Rh, \dots, nRh$ . We denote with  $N$  the number of elements in the discretisation  $\{\tau_1, \tau_h, \dots, \tau_{Rh}, \dots, \tau_{nRh}\}$  and with  $\mathbf{X}_{1:N} = (\mathbf{X}_{\tau_1}, \dots, \mathbf{X}_{\tau_N})$  the corresponding values of  $\{\mathbf{X}_t\}$  obtained when using a given numerical/approximated method of choice. Then the likelihood function becomes

$$\begin{aligned} p(\mathbf{Y}_{1:n}; \boldsymbol{\theta}) &= \int p_{\mathbf{Y}, \mathbf{X}}(\mathbf{Y}_{1:n}, \mathbf{X}_{0:N}; \boldsymbol{\theta}) d\mathbf{X}_{0:N} = \int p_{\mathbf{Y}|\mathbf{X}}(\mathbf{Y}_{1:n}|\mathbf{X}_{0:N}; \boldsymbol{\theta}) p_{\mathbf{X}}(\mathbf{X}_{0:N}; \boldsymbol{\theta}) d\mathbf{X}_{0:N} \\ &= \int \left\{ \prod_{j=1}^n f(\mathbf{Y}_j|\mathbf{X}_j; \boldsymbol{\theta}) \right\} p(\mathbf{X}_0) \prod_{i=1}^N p(\mathbf{X}_i|\mathbf{X}_{i-1}; \boldsymbol{\theta}) d\mathbf{X}_0 \cdots d\mathbf{X}_N, \end{aligned}$$

where the product in  $j$  is over the  $\mathbf{X}_{t_j}$  and the product in  $i$  is over the  $\mathbf{X}_{\tau_i}$ .

### 2.1 The standard SAEM algorithm

The EM algorithm introduced by [Dempster et al. \[1977\]](#) is a classical approach to estimate parameters by maximum likelihood for models with non-observed or incomplete data. Let us briefly cover the EM principle. The complete data of the model is  $(\mathbf{Y}_{1:n}, \mathbf{X}_{0:N})$ , where  $\mathbf{X}_{0:N} \equiv \mathbf{X}_{0:n}$  if numerical discretization is not required, and for ease of writing we denote this as  $(\mathbf{Y}, \mathbf{X})$  in the remaining of this section. The EM algorithm maximizes the function  $Q(\boldsymbol{\theta}|\boldsymbol{\theta}') = \mathbb{E}(L_c(\mathbf{Y}, \mathbf{X}; \boldsymbol{\theta})|\mathbf{Y}; \boldsymbol{\theta}')$  in two steps, where  $L_c(\mathbf{Y}, \mathbf{X}; \boldsymbol{\theta}) := \log p_{\mathbf{Y}, \mathbf{X}}(\mathbf{Y}, \mathbf{X}; \boldsymbol{\theta})$  is the log-likelihood of the complete data and  $\mathbb{E}$  is the conditional expectation under the conditional distribution  $p_{\mathbf{X}|\mathbf{Y}}(\cdot; \boldsymbol{\theta}')$ . At the  $k$ -th iteration of a maximum (user defined) number of evaluations  $K$ , the E-step is the evaluation of  $Q_k(\boldsymbol{\theta}) = Q(\boldsymbol{\theta}|\hat{\boldsymbol{\theta}}^{(k-1)})$ , whereas the M-step updates  $\hat{\boldsymbol{\theta}}^{(k-1)}$  by maximizing  $Q_k(\boldsymbol{\theta})$ . For cases in which the E-step has no analytic form, [Delyon et al. \[1999\]](#) introduce a stochastic version of the EM algorithm (SAEM) which evaluates the integral  $Q_k(\boldsymbol{\theta})$  by a stochastic approximation procedure. The authors prove the convergence of this algorithm under general conditions if  $L_c(\mathbf{Y}, \mathbf{X}; \boldsymbol{\theta})$  belongs to the regular exponential family

$$L_c(\mathbf{Y}, \mathbf{X}; \boldsymbol{\theta}) = -\Lambda(\boldsymbol{\theta}) + \langle \mathbf{S}_c(\mathbf{Y}, \mathbf{X}), \boldsymbol{\Gamma}(\boldsymbol{\theta}) \rangle,$$

where  $\langle \cdot, \cdot \rangle$  is the scalar product,  $\Lambda$  and  $\Gamma$  are two functions of  $\boldsymbol{\theta}$  and  $\mathbf{S}_c(\mathbf{Y}, \mathbf{X})$  is the minimal sufficient statistic of the complete model. The E-step is then divided into a simulation step (S-step) of the missing data  $\mathbf{X}^{(k)}$  under the conditional distribution  $p_{\mathbf{X}|\mathbf{Y}}(\cdot; \hat{\boldsymbol{\theta}}^{(k-1)})$  and a stochastic approximation step (SA-step) of the conditional expectation, using  $(\gamma_k)_{k \geq 1}$  a sequence of real numbers in  $[0, 1]$ , such that  $\sum_{k=1}^{\infty} \gamma_k = \infty$  and  $\sum_{k=1}^{\infty} \gamma_k^2 < \infty$ . This SA-step approximates  $\mathbb{E} \left[ \mathbf{S}_c(\mathbf{Y}, \mathbf{X}) | \hat{\boldsymbol{\theta}}^{(k-1)} \right]$  at each iteration by the value  $\mathbf{s}_k$  defined recursively as follows

$$\mathbf{s}_k = \mathbf{s}_{k-1} + \gamma_k (\mathbf{S}_c(\mathbf{Y}, \mathbf{X}^{(k)}) - \mathbf{s}_{k-1}). \quad (3)$$

The M-step is thus the update of the estimates  $\hat{\boldsymbol{\theta}}^{(k-1)}$

$$\hat{\boldsymbol{\theta}}^{(k)} = \arg \max_{\boldsymbol{\theta} \in \Theta} (-\Lambda(\boldsymbol{\theta}) + \langle \mathbf{s}_k, \Gamma(\boldsymbol{\theta}) \rangle). \quad (4)$$

The starting  $\mathbf{s}_0$  can be set to be a vector of zeros. The procedure above can be carried out iteratively for  $K$  iterations. The proof of the convergence of SAEM requires the sequence  $\gamma_k$  to be such that  $\sum_{k=1}^{\infty} \gamma_k = \infty$  and  $\sum_{k=1}^{\infty} \gamma_k^2 < \infty$ . A typical choice is to consider a warmup period with  $\gamma_k = 1$  for the first  $K_1$  iterations and then  $\gamma_k = (k - K_1)^{-1}$  for  $k \geq K_1$  (with  $K_1 < K$ ). Parameter  $K_1$  has to be chosen by the user. However inference results are typically not very sensitive to this tuning parameter. Typical values are  $K_1 = 250$  or  $300$  and  $K = 400$ , see for example [Lavielle \[2014\]](#). Usually, the simulation step of the hidden trajectory  $\mathbf{X}^{(k)}$  conditionally on the observations  $\mathbf{Y}$  cannot be directly performed. [Lindsten \[2013\]](#) and [Ditlevsen and Samson \[2014\]](#) use a sequential Monte Carlo algorithm to perform the simulation step for state-space models. We propose to resort to approximate Bayesian computation (ABC) for this simulation step.

It can be noted that the generation of  $\mathbf{s}_k$  in (3) followed by corresponding parameter estimates  $\hat{\boldsymbol{\theta}}^{(k)}$  in (4) is akin to two steps of a Gibbs sampling algorithm, except that here  $\hat{\boldsymbol{\theta}}^{(k)}$  is produced by a deterministic step (for given  $\mathbf{s}_k$ ). We comment further on this aspect in section 4.1.4.

## 2.2 The SAEM algorithm coupled to an ABC simulation step

Approximate Bayesian Computation (ABC, see [Marin et al., 2012](#) for a review) is a class of probabilistic algorithms allowing sampling from an approximation of a posterior distribution. The most typical use of ABC is when posterior inference for  $\boldsymbol{\theta}$  is the goal of the analysis and the purpose is to sample draws from the approximate posterior  $\pi_{\delta}(\boldsymbol{\theta}|\mathbf{Y})$ . Here and in the following  $\mathbf{Y} \equiv \mathbf{Y}_{1:n}$ . The parameter  $\delta > 0$  is a threshold influencing the quality of the inference, the smaller the  $\delta$  the more accurate the inference, and  $\pi_{\delta}(\boldsymbol{\theta}|\mathbf{Y}) \equiv \pi(\boldsymbol{\theta}|\mathbf{Y})$  when  $\delta = 0$ . However in our study we are not interested in conducting Bayesian inference for  $\boldsymbol{\theta}$ . We will use ABC to sample from an approximation to the posterior distribution  $\pi(\mathbf{X}_{0:N}|\mathbf{Y}; \boldsymbol{\theta}) \equiv p(\mathbf{X}_{0:N}|\mathbf{Y}; \boldsymbol{\theta})$ , that is for a fixed value of  $\boldsymbol{\theta}$ , we wish to sample from  $\pi_{\delta}(\mathbf{X}_{0:N}|\mathbf{Y}; \boldsymbol{\theta})$  (recall from section 2.1 that when feasible we can take  $N \equiv n$ ). For simplicity of notation, in the following we avoid specifying the dependence on the current value of  $\boldsymbol{\theta}$ , which has to be assumed as a deterministic unknown. There are several ways to generate a ‘‘candidate’’  $\mathbf{X}_{0:N}^*$ : for example we might consider ‘‘blind’’ forward simulation, meaning that  $\mathbf{X}_{0:N}^*$  is simulated from  $p_{\mathbf{X}}(\mathbf{X}_{0:N})$  and therefore unconditionally on data (i.e. the simulator is blind with respect to data). Then  $\mathbf{X}_{0:N}^*$  is accepted if the corresponding  $\mathbf{Y}^*$  simulated from  $f(\cdot|\mathbf{X}_{1:n}^*)$  is ‘‘close’’ to  $\mathbf{Y}$ , according to the threshold  $\delta$ , where  $\mathbf{X}_{1:n}^*$  contains the interpolated values of  $\mathbf{X}_{0:N}^*$  at sampling times  $\{t_1, \dots, t_n\}$  and  $\mathbf{Y}^* \equiv \mathbf{Y}_{1:n}^*$ . Notice that the appeal of the methodology is that knowledge of the probabilistic features of the data generating model is not necessary, meaning that even if the transition densities  $p(\mathbf{X}_j|\mathbf{X}_{j-1})$  are not known (hence  $p_{\mathbf{X}}$  is unknown) it is enough to be able to simulate from the model (using a numerical scheme if necessary) hence draws  $\mathbf{X}_{0:N}^*$  are produced by forward-simulation regardless the explicit knowledge of the underlying densities.

Algorithm 1 illustrates a generic iteration  $k$  of a SAEM-ABC method, where the current value of the parameters is  $\hat{\boldsymbol{\theta}}^{(k-1)}$  and an updated value of the estimates is produced as  $\hat{\boldsymbol{\theta}}^{(k)}$ . By iterating the procedure  $K$  times, the resulting  $\hat{\boldsymbol{\theta}}^{(k)}$  is the maximizer for an approximate

---

**Algorithm 1** A generic iteration of SAEM-ABC using acceptance-rejection

---

**Simulation step:** here we update  $\mathbf{X}^{(k)}$  using an ABC procedure sampling from  $\pi_{\delta}(\mathbf{X}|\mathbf{Y}; \hat{\boldsymbol{\theta}}^{(k-1)})$ :

**Repeat**

- Generate a candidate  $\mathbf{X}^*$  from the latent model dynamics conditionally on  $\hat{\boldsymbol{\theta}}^{(k-1)}$ , either by numerical methods or using the transition density (if available) i.e. by generating using the exact law  $p_{\mathbf{X}}(\cdot; \hat{\boldsymbol{\theta}}^{(k-1)})$ .
- Generate  $\mathbf{Y}^*$  from the error model  $f(\mathbf{Y}^*|\mathbf{X}^*)$ .

**Until**  $\rho(\mathbf{Y}^*, \mathbf{Y}) \leq \delta$

Set  $\mathbf{X}^{(k)} = \mathbf{X}^*$

**Stochastic Approximation step :** update of the sufficient statistics

$$\mathbf{s}_k = \mathbf{s}_{k-1} + \gamma_k \left( \mathbf{S}_c(\mathbf{Y}, \mathbf{X}^{(k)}) - \mathbf{s}_{k-1} \right)$$

**Maximisation step:** update  $\boldsymbol{\theta}$

$$\hat{\boldsymbol{\theta}}^{(k)} = \arg \max_{\boldsymbol{\theta} \in \Theta} (-\Lambda(\boldsymbol{\theta}) + \langle \mathbf{s}_k, \Gamma(\boldsymbol{\theta}) \rangle)$$


---

likelihood (the approximation implied by using ABC). The “repeat loop” can be considerably expensive using a distance  $\rho(\mathbf{Y}^*, \mathbf{Y}) \leq \delta$  as acceptance of  $\mathbf{Y}^*$  (hence acceptance of  $\mathbf{X}^*$ ) is a rare event for  $\delta$  reasonably small. If informative low-dimensional statistics  $\eta(\cdot)$  are available, it is recommended to consider  $\rho(\eta(\mathbf{Y}^*), \eta(\mathbf{Y}))$  instead.

However algorithm 1 is not appropriate for state-space models, because the entire candidate trajectory  $\mathbf{X}^*$  is simulated blindly to data (an alternative approach is considered in section 3). If we consider for a moment  $\mathbf{X}^* \equiv \mathbf{X}_{0:N}^*$  as a generic unknown, ideally we would like to sample from the posterior  $\pi(\mathbf{X}^*|\mathbf{Y})$ , which is proportional to

$$\pi(\mathbf{X}^*|\mathbf{Y}) \propto f(\mathbf{Y}|\mathbf{X}^*)\pi(\mathbf{X}^*) \tag{5}$$

for a given “prior” distribution  $\pi(\mathbf{X}^*)$ . For some models sampling from such posterior is not trivial, for example when  $\mathbf{X}$  is a stochastic process and in that case sequential Monte Carlo methods can be used as described in section 3. A further layer of approximation is introduced when  $\pi(\mathbf{X}^*|\mathbf{Y})$  is analytically “intractable”, and specifically when  $f(\mathbf{Y}|\mathbf{X})$  is unavailable in closed form (though we always assume  $f(\cdot|\cdot)$  to be known and that it is possible to evaluate it pointwise) or is computationally difficult to approximate but it is easy to sample from. In this case ABC methodology turns useful, and an ABC approximation to  $\pi(\mathbf{X}^*|\mathbf{Y})$  can be written as

$$\pi_{\delta}(\mathbf{X}^*|\mathbf{Y}) \propto \int J_{\delta}(\mathbf{Y}, \mathbf{Y}^*) \underbrace{f(\mathbf{Y}^*|\mathbf{X}^*)\pi(\mathbf{X}^*)}_{\propto \pi(\mathbf{X}^*|\mathbf{Y}^*)} d\mathbf{Y}^*. \tag{6}$$

Here  $J_{\delta}(\cdot)$  is some function that depends on  $\delta$  and weights the intractable posterior based on simulated data  $\pi(\mathbf{X}^*|\mathbf{Y}^*)$  with high values in regions where  $\mathbf{Y}$  and  $\mathbf{Y}^*$  are similar; therefore we would like (i)  $J_{\delta}(\cdot)$  to give higher rewards to proposals  $\mathbf{X}^*$  corresponding to  $\mathbf{Y}^*$  having values close to  $\mathbf{Y}$ . In addition (ii)  $J_{\delta}(\mathbf{Y}, \mathbf{Y}^*)$  is assumed to be a constant when  $\mathbf{Y}^* = \mathbf{Y}$  (i.e. when

$\delta = 0$ ) so that  $J_\delta$  is absorbed into the proportionality constant and the *exact* marginal posterior  $\pi(\mathbf{X}^*|\mathbf{Y})$  is recovered. Basically the use of (6) means to simulate  $\mathbf{X}^*$  from its prior (the product of transition densities), then plug such draw into  $f(\cdot|\mathbf{X}^*)$  to simulate  $\mathbf{Y}^*$ , so that  $\mathbf{X}^*$  will be weighted by  $J_\delta(\mathbf{Y}, \mathbf{Y}^*)$ . A common choice for  $J_\delta(\cdot)$  is the uniform kernel

$$J_\delta(\mathbf{Y}, \mathbf{Y}^*) \propto \mathbb{I}_{\{\rho(\mathbf{Y}^*, \mathbf{Y}) \leq \delta\}}$$

where  $\rho(\mathbf{Y}, \mathbf{Y}^*)$  is some measure of closeness between  $\mathbf{Y}^*$  and  $\mathbf{Y}$  and  $\mathbb{I}$  is the indicator function. A further popular possibility is the Gaussian kernel

$$J_\delta(\mathbf{Y}, \mathbf{Y}^*) \propto e^{-(\mathbf{Y}^* - \mathbf{Y})'(\mathbf{Y}^* - \mathbf{Y})/2\delta^2} \quad (7)$$

where  $'$  denotes transposition, so that  $J_\delta(\mathbf{Y}, \mathbf{Y}^*)$  gets larger when  $\mathbf{Y}^* \approx \mathbf{Y}$ .

However one of the difficulties is that, in practice,  $\delta$  has to be set as a compromise between statistical accuracy (a small positive  $\delta$ ) and computational feasibility ( $\delta$  not too small). Notice that a proposal's acceptance can be significantly enhanced when the posterior (6) is conditional on summary statistics of data  $\eta(\mathbf{Y})$ , rather than  $\mathbf{Y}$  itself, and in such case we would consider  $\rho(\eta(\mathbf{Y}), \eta(\mathbf{Y}^*))$ . However, in practice for dynamical models it is difficult to identify “informative enough” summary statistics  $\eta(\cdot)$ , but see [Martin et al. \[2016\]](#) and [Picchini and Forman \[2015\]](#). Another important problem with the strategy outlined above is that “blind simulation” for the generation of the entire time series  $\mathbf{X}^*$  is often poor. In fact, even when the current value of  $\theta$  is close to its true value, proposed trajectories rarely follow measurements when (a) the dataset is a long time series and/or (b) the model is highly erratic, for example when latent dynamics are expressed by a stochastic differential equation (section 4.2). For these reasons, sequential Monte Carlo (SMC) [[Cappe et al., 2007](#)] methods have emerged as the most successful solution for filtering in non-linear non-Gaussian state-space models.

Below we consider the ABC-SMC methodology from [Jasra et al. \[2012\]](#), which proves effective for state-space models.

### 3 SAEM coupled with an ABC-SMC algorithm for filtering

Here we consider a strategy for filtering that is based on an ABC version of sequential Monte Carlo sampling, as presented in [Jasra et al. \[2012\]](#), with some minor modifications. The advantage of the methodology is that the generation of proposed trajectories is sequential, that is the ABC distance is evaluated “locally” for each observational time point. In fact, what is evaluated is the proximity of trajectories/particles to each data point  $\mathbf{Y}_j$ , and “bad trajectories” are killed thus preventing the propagation of unlikely states to the next observation  $\mathbf{Y}_{j+1}$  and so on. For simplicity we consider the case  $N \equiv n$ ,  $h \equiv \Delta$ . The algorithm samples from the following target density at time  $t_{n'}$  ( $n' \leq n$ ):

$$\pi_{n', \delta}(\mathbf{X}_{1:n'}, \mathbf{Y}_{1:n'}^* | \mathbf{Y}_{1:n'}) \propto \prod_{j=1}^{n'} J_{j, \delta}(\mathbf{Y}_j, \mathbf{Y}_j^*) f(\mathbf{Y}_j^* | \mathbf{X}_j) p(\mathbf{X}_j | \mathbf{X}_{j-1})$$

and for example we could take  $J_{j, \delta}(\mathbf{y}_j, \mathbf{y}_j^*) = \mathbb{I}_{A_{\delta, \mathbf{y}_j}}(\mathbf{y}_j^*)$  with  $A_{\delta, \mathbf{y}_j} = \{\mathbf{y}_j^*; \rho(\eta(\mathbf{y}_j^*), \eta(\mathbf{y}_j)) < \delta\}$  as in [Jasra et al. \[2012\]](#) or the Gaussian kernel (7).

The ABC-SMC procedure is set in algorithm 2 with the purpose to propagate forward  $M$  simulated states (“particles”). After algorithm 2 is executed, we select a single trajectory by retrospectively looking at the genealogy of the generated particles, as explained further below. The quantity ESS is the effective sample size (e.g. [Liu, 2008](#)) often estimated as  $ESS(\{w_j^{(m)}\}) = 1 / \sum_{m=1}^M (w_j^{(m)})^2$  and taking values between 1 and  $M$ . When considering an indicator function for  $J_{j, \delta}$ , the ESS coincides with the number of particles having positive weight [[Jasra et al., 2012](#)]. Under such choice the integer  $\bar{M} \leq M$  is a lower bound (threshold set by the experimenter) on

---

**Algorithm 2** ABC-SMC for filtering
 

---

Step 0. Set  $j = 1$ . For  $m = 1, \dots, M$  sample  $\mathbf{X}_1^{(m)} \sim p(\mathbf{X}_0)$ ,  $\mathbf{Y}_1^{*(m)} \sim f(\cdot | \mathbf{X}_1^{(m)})$ , compute weights  $W_1^{(m)} = J_{1,\delta}(\mathbf{Y}_1, \mathbf{Y}_1^{*(m)})$  and normalize weights  $w_1^{(m)} := W_1^{(m)} / \sum_{m=1}^M W_1^{(m)}$ .

Step 1.

**if**  $ESS(\{w_j^{(m)}\}) < \bar{M}$  **then**

resample  $M$  particles  $\{\mathbf{X}_j^{(m)}, w_j^{(m)}\}$  and set  $W_j^{(m)} = 1/M$ .

**end if**

Set  $j := j + 1$  and if  $j = n + 1$ , stop.

Step 2. For  $m = 1, \dots, M$  sample  $\mathbf{X}_j^{(m)} \sim p(\cdot | \mathbf{X}_{j-1}^{(m)})$  and  $\mathbf{Y}_j^{*(m)} \sim f(\cdot | \mathbf{X}_j^{(m)})$ . Compute

$$W_j^{(m)} := w_{j-1}^{(m)} J_{j,\delta}(\mathbf{Y}_j, \mathbf{Y}_j^{*(m)})$$

normalize weights  $w_j^{(m)} := W_j^{(m)} / \sum_{m=1}^M W_j^{(m)}$  and go to step 1.

---

the number of particles with non-zero weight. However in our experiments we use a Gaussian kernel and since in the examples in section 4 we have a scalar  $Y_j$ , the kernel is defined as

$$J_{j,\delta}(Y_j, Y_j^{*(m)}) = \frac{1}{\delta} e^{-(Y_j^{*(m)} - Y_j)^2 / (2\delta^2)} \quad (8)$$

so that weights  $W_j^{(m)}$  are larger for particles having  $Y_j^{*(m)} \approx Y_j$ . We consider “stratified resampling” [Kitagawa, 1996] in step 1 of algorithm 2.

In addition to the procedure outlined in algorithm 2, once the set of weights  $\{w_n^{(1)}, \dots, w_n^{(M)}\}$  is available at time  $t_n$ , we propose to follow Andrieu et al. [2010] (see their PMMH algorithm) and sample a single index from the set  $\{1, \dots, M\}$  having associated probabilities  $\{w_n^{(1)}, \dots, w_n^{(M)}\}$ . Denote with  $m'$  such index and with  $a_j^m$  the “ancestor” of the generic  $m$ th particle sampled at time  $t_{j+1}$ , with  $1 \leq a_j^m \leq M$  ( $m = 1, \dots, M$ ,  $j = 1, \dots, n$ ). Then we have that particle  $m'$  has ancestor  $a_{n-1}^{m'}$  and in general particle  $m''$  at time  $t_{j+1}$  has ancestor  $b_j^{m''} := a_j^{b_{j+1}^{m''}}$ , with  $b_n^{m'} := m'$ . Hence, at the end of algorithm 2 we can sample  $m'$  and construct its genealogy: the sequence of states  $\{X_t\}$  resulting from the genealogy of  $m'$  is the chosen path that will be passed to SAEM, see algorithm 3. The selection of this path is crucially affected by “particles impoverishment” issues, see below. [An alternative procedure, which we do not pursue here, is to sample at time  $t_n$  not just a single index  $m'$ , but instead sample with replacement say  $G \geq 1$  times from the set  $\{1, \dots, M\}$  having associated probabilities  $\{w_n^{(1)}, \dots, w_n^{(M)}\}$ . Then construct the genealogy for each of the  $G$  sampled indices, and for each of the resulting  $G$  sampled paths  $\mathbf{X}^{g,k}$  calculate the corresponding vector-summaries  $\mathbf{S}_c^{g,k} := \mathbf{S}_c(\mathbf{Y}, \mathbf{X}^{g,k})$ ,  $g = 1, \dots, G$ . Then it would be possible to take the sample average of those  $G$  summaries  $\bar{\mathbf{S}}_c^k$  as in Kuhn and Lavielle [2005], and plug this average in place of  $\mathbf{S}_c(\mathbf{Y}, \mathbf{X}^{(k)})$  in step 3 of algorithm 3. Clearly this approach increases the computational complexity linearly with  $G$ .]

Notice that in Jasra et al. [2012]  $n$  ABC thresholds  $\{\delta_1, \dots, \delta_n\}$  are constructed, one threshold for each corresponding sampling time in  $\{t_1, \dots, t_n\}$ : these thresholds do not need to be set by the user but can be updated adaptively using a stochastic data-driven procedure. This is possible because the ABC-SMC algorithm in Jasra et al. [2012] is for filtering only, that is  $\theta$  is a fixed known quantity. However in our scenario  $\theta$  is unknown, and letting the thresholds vary adaptively (and randomly) between each pair of iterations of a parameter estimation algorithm is not appropriate. This is because the evaluation of the likelihood function at two different iterations  $k'$  and  $k''$  of SAEM would depend on a procedure determining corresponding (stochastic) sequences  $\{\delta_1^{(k')}, \dots, \delta_n^{(k')}\}$  and  $\{\delta_1^{(k'')}, \dots, \delta_n^{(k'')}\}$ . Therefore the likelihood maximization would be affected by the random realizations of the sequences of thresholds. In our case, we let the threshold vary



*deterministically*: that is we choose a  $\delta$  common to all time points  $\{t_1, \dots, t_n\}$  and execute a number of SAEM iterations using such threshold. Then we deterministically decrease the threshold according to a user-defined scheme and execute further SAEM iterations, and so on. Our SAEM-ABC procedure is detailed in algorithm 3 with a user-defined sequence  $\delta_1 > \dots > \delta_L > 0$  where each  $\delta_l$  is used for  $k_l$  iterations, so that  $k_1 + \dots + k_L = K$  ( $l = 1, \dots, L$ ). Here  $K$  is the number of SAEM iterations, as defined at the end of section 2.1. In our applications we show how the algorithm is not overly sensitive to small variations in the  $\delta$ 's.

Regarding the choice of  $\delta$  values in applications, recall that the interpretation of  $\delta$  in equation (8) is that of the standard deviation for a perturbed model. This implies that a synthetic observation at time  $t_j$ , denoted with  $Y_j^*$ , can be interpreted as a perturbed version of  $Y_j$ , where the observed  $Y_j$  is assumed generated from the state-space model in equation (1), while  $Y_j^* \sim \mathcal{N}(Y_j, \delta^2)$ . With this fact in mind,  $\delta$  can easily be chosen to represent some deviation from the actual observations. Therefore, typically it is enough to look at the time evolution of the data, to guess at least the order of magnitude of the starting value for  $\delta$ .

Finally, in our applications we compare SAEM-ABC with SAEM-SMC. SAEM-SMC is detailed in algorithm 4: it is structurally the same as algorithm 3 except that it uses the bootstrap filter to select the state trajectory. The bootstrap filter [Gordon et al., 1993] is just like algorithm 2, except that no simulation from the observations equation is performed (i.e. the  $\mathbf{Y}_j^{*(m)}$  are not generated) and  $J_{j,\delta}(\mathbf{Y}_j, \mathbf{Y}_j^{*(m)})$  is replaced with  $f(\mathbf{Y}_j | \mathbf{X}_j^{(m)})$ , hence there is no need to specify the  $\delta$ . The trajectory  $\mathbf{X}^{(k)}$  selected by the bootstrap filter at  $k$ th SAEM iteration is then used to update the statistics  $\mathbf{s}_k$ .

As studied in detail in section 4.3, the function  $J_{t,\delta}$  has an important role in approximating the density  $\pi(\mathbf{X}_{1:t} | \mathbf{Y}_{1:t})$  at a generic time instant  $t$ . Indeed  $\pi_\delta(\mathbf{X}_{1:t} | \mathbf{Y}_{1:t}) = \int \pi_\delta(\mathbf{X}_{1:t}, \mathbf{Y}_{1:t}^* | \mathbf{Y}_{1:t}) d\mathbf{Y}_{1:t}^*$ , the ABC approximation to  $\pi(\mathbf{X}_{1:t} | \mathbf{Y}_{1:t})$ , is

$$\pi_\delta(\mathbf{X}_{1:t} | \mathbf{Y}_{1:t}) = \sum_{m=1}^M W_t^{(m)} \xi_{\mathbf{x}_{0:t}}(\mathbf{x}_{0:t}^{(m)})$$

with  $\xi(\cdot)$  the Dirac measure and  $W_t^{(m)} \propto J_{t,\delta}(\mathbf{Y}_t^{*(m)}, \mathbf{Y}_t)$ . This means that for a small  $\delta$  (which is set by the user)  $J_\delta$  assigns large weights only to very promising particles, that is those particles  $\mathbf{X}_t^{(m)}$  associated to a  $\mathbf{Y}_t^{*(m)}$  very close to  $\mathbf{Y}_t$ . Instead, suppose that  $f(\cdot | \mathbf{X}_t, \boldsymbol{\theta}_y) \equiv \mathcal{N}(\cdot | \mathbf{X}_t; \sigma_\varepsilon^2)$  where  $\mathcal{N}(\cdot | a; b)$  is a Gaussian distribution with mean  $a$  and variance  $b$ , then the bootstrap filter underlying SAEM-SMC assigns weights proportionally to the measurements density  $f(\mathbf{Y}_t | \mathbf{X}_t; \sigma_\varepsilon)$ , which is affected by the currently available (and possibly poor) value of  $\sigma_\varepsilon$ , when  $\sigma_\varepsilon$  is one of the unknowns to be estimated. When  $\sigma_\varepsilon$  is overestimated, as in section 4.3, there is a risk to sample particles which are not really ‘‘important’’. Since trajectories selected in step 2 of algorithm 4 are drawn from  $\pi(\mathbf{X}_{1:t} | \mathbf{Y}_{1:t})$ , clearly the issues just highlighted contribute to bias the inference.

A further issue is studied in section 4.1.3. There we explain how SAEM-ABC, despite being an approximate version of SAEM-SMC (due to the additional approximation induced by using a strictly positive  $\delta$ ) can in practice outperform SAEM coupled with the simple bootstrap filter, because of ‘‘particles impoverishment’’ problems. Particles impoverishment is a pathology due to frequently implementing particles resampling: the resampling step reduces the ‘‘variety’’ of the particles, by duplicating the ones with larger weights and killing the others. We show that when particles fail to get close to the targeted observations, then resampling is frequently triggered, this degrading the variety of the particles. However with an ABC filter, particles receive some additional weighting, due to the function  $J_{j,\delta}$  in equation (8) taking values  $J_{j,\delta} > 1$  for small  $\delta$  when  $\mathbf{Y}_j^* \approx \mathbf{Y}_j$ , which allows a larger number of particles to have a non-negligible weight and therefore different particles are resampled, this increasing their variety (at least for the application in section 4.1.1, but this is not true in general). This is especially relevant in early iterations of SAEM, where  $\boldsymbol{\theta}^{(k)}$  might be far from its true value and therefore many particles

---

**Algorithm 3** SAEM-ABC using a particle filter

---

Step 0. Set parameters starting values  $\hat{\boldsymbol{\theta}}^{(0)}$ , set  $M, \bar{M}$  and  $k := 1$ . Set the sequence  $\{\delta_1, \dots, \delta_L\}$  and  $\delta := \delta_1$ .

Step 1. For fixed  $\hat{\boldsymbol{\theta}}^{(k-1)}$  apply the ABC-SMC algorithm 2 with threshold  $\delta$ ,  $M$  particles and particles threshold  $\bar{M}$ .

2 Sample an index  $m'$  from the probability distribution  $\{w_n^{(1)}, \dots, w_n^{(M)}\}$  on  $\{1, \dots, M\}$  and form the path  $\mathbf{X}^{(k)}$  resulting from the genealogy of  $m'$ .

Step 3. **Stochastic Approximation step** : update of the sufficient statistics

$$\mathbf{s}_k = \mathbf{s}_{k-1} + \gamma_k \left( \mathbf{S}_c(\mathbf{Y}, \mathbf{X}^{(k)}) - \mathbf{s}_{k-1} \right)$$

Step 4. **Maximisation step**: update  $\boldsymbol{\theta}$

$$\hat{\boldsymbol{\theta}}^{(k)} = \arg \max_{\boldsymbol{\theta} \in \Theta} (-\Lambda(\boldsymbol{\theta}) + \langle \mathbf{s}_k, \Gamma(\boldsymbol{\theta}) \rangle)$$

Set  $k := k + 1$ . If  $k \in \{k_1, \dots, k_L\}$ , e.g. if  $k = k_L$ , set  $\delta := \delta_L$ . Go to step 1.

---

---

**Algorithm 4** SAEM-SMC

---

Step 0. Set parameters starting values  $\hat{\boldsymbol{\theta}}^{(0)}$ , set  $M, \bar{M}$  and  $k := 1$ . For fixed  $\hat{\boldsymbol{\theta}}^{(k-1)}$  apply the bootstrap filter below using  $M$  particles and particles threshold  $\bar{M}$ .

Step 1a. Set  $j = 1$ . For  $m = 1, \dots, M$  and conditionally on  $\hat{\boldsymbol{\theta}}^{(k-1)}$  sample  $\mathbf{X}_1^{(m)} \sim p(\mathbf{X}_0)$ , compute weights  $W_1^{(m)} = f(\mathbf{Y}_1 | \mathbf{X}_1^{(m)})$  and normalize weights  $w_1^{(m)} := W_1^{(m)} / \sum_{m=1}^M W_1^{(m)}$ .

Step 1.b.

**if**  $ESS(\{w_j^{(m)}\}) < \bar{M}$  **then** resample  $M$  particles  $\{\mathbf{X}_j^{(m)}, w_j^{(m)}\}$  and set  $W_j^{(m)} = 1/M$ .

**end if**

Set  $j := j + 1$ . If  $j = n + 1$  go to step 2. If  $j \leq n$  go to step 1.c.

Step 1.c. For  $m = 1, \dots, M$  and conditionally on  $\hat{\boldsymbol{\theta}}^{(k-1)}$  sample  $\mathbf{X}_j^{(m)} \sim p(\cdot | \mathbf{X}_{j-1}^{(m)})$ . Compute

$$W_j^{(m)} := w_{j-1}^{(m)} f(\mathbf{Y}_j | \mathbf{X}_j^{(m)})$$

normalize weights  $w_j^{(m)} := W_j^{(m)} / \sum_{m=1}^M W_j^{(m)}$  and go to step 1.b.

Step 2. Sample an index  $m'$  from the probability distribution  $\{w_n^{(1)}, \dots, w_n^{(M)}\}$  on  $\{1, \dots, M\}$  and form the path  $\mathbf{X}^{(k)}$  resulting from the genealogy of  $m'$ .

Step 3. **Stochastic Approximation step** : update of the sufficient statistics

$$\mathbf{s}_k = \mathbf{s}_{k-1} + \gamma_k \left( \mathbf{S}_c(\mathbf{Y}, \mathbf{X}^{(k)}) - \mathbf{s}_{k-1} \right)$$

Step 4. **Maximisation step**: update  $\boldsymbol{\theta}$

$$\hat{\boldsymbol{\theta}}^{(k)} = \arg \max_{\boldsymbol{\theta} \in \Theta} (-\Lambda(\boldsymbol{\theta}) + \langle \mathbf{s}_k, \Gamma(\boldsymbol{\theta}) \rangle)$$

Set  $k := k + 1$ . Go to step 1a.

---

might end far from data. By letting  $\delta$  decrease not “too fast” as SAEM iterations increase, we allow many particles to contribute to the states propagation. However, as  $\delta$  approaches a small value, only the most promising particles will contribute to selecting the path sampled in step 2 of algorithm 3.

These aspects are discussed in greater detail in sections 4.1.3 and 4.3. However, it is of course *not true* that an ABC filter is in general expected to perform better than a non-ABC one. In the second example, section 4.2, an adapted (not “blind to data”, i.e. conditional to the next observation) particle filter is used to treat the specific case of SDE models requiring numerical integration [Golightly and Wilkinson, 2011]. The adapted filter clearly outperforms both SAEM-SMC and SAEM-ABC when the number of particles is very limited ( $M = 30$ ).

Notice, our SAEM-ABC strategy with a decreasing series of thresholds shares some similarity with tempering approaches. For example Herbst and Schorfheide [2017] artificially inflate the observational noise variance pertaining  $f(\mathbf{y}_t|\mathbf{X}_t, \cdot)$ , so that the particle weights have lower variance hence the resulting filter is more stable. More in detail, they construct a bootstrap filter where particles are propagated through a sequence of intermediate tempering steps, starting from an observational distribution with inflated variance, and then gradually reducing the variance to its nominal level.

**Fisher Information matrix** The SAEM algorithm allows also to compute standard errors of the estimators, through the approximation of the Fisher Information matrix. This is detailed below, however notice that the algorithm itself advances between iterations without the need to compute such matrix (nor the gradient of the function to maximize). The standard errors of the parameter estimates can be calculated from the diagonal elements of the inverse of the Fisher information matrix. Its direct evaluation is difficult because it has no explicit analytic form, however an estimate of the Fisher information matrix can easily be implemented within SAEM as proposed by Delyon et al. [1999] using the Louis’ missing information principle [Louis, 1982].

The Fisher information matrix  $\ell(\boldsymbol{\theta}) = L(\mathbf{Y}; \boldsymbol{\theta})$  can be expressed as:

$$\begin{aligned} \partial_{\theta}^2 \ell(\boldsymbol{\theta}) &= \mathbb{E} [\partial_{\theta}^2 L_c(\mathbf{Y}, \mathbf{X}; \boldsymbol{\theta}) | \mathbf{Y}, \boldsymbol{\theta}] \\ &+ \mathbb{E} [\partial_{\theta} L_c(\mathbf{Y}, \mathbf{X}; \boldsymbol{\theta}) (\partial_{\theta} L_c(\mathbf{Y}, \mathbf{X}; \boldsymbol{\theta}))' | \mathbf{Y}, \boldsymbol{\theta}] \\ &- \mathbb{E} [\partial_{\theta} L_c(\mathbf{Y}, \mathbf{X}; \boldsymbol{\theta}) | \mathbf{Y}, \boldsymbol{\theta}] \mathbb{E} [\partial_{\theta} L_c(\mathbf{Y}, \mathbf{X}; \boldsymbol{\theta}) | \mathbf{Y}, \boldsymbol{\theta}]' \end{aligned}$$

where  $'$  denotes transposition. An on-line estimation of the Fisher information is obtained using the stochastic approximation procedure of the SAEM algorithm as follows (see Lavielle, 2014 for an off-line approach). At the  $(k+1)$ th iteration of the algorithm, we evaluate the three following quantities:

$$\begin{aligned} \mathbf{G}_{k+1} &= \mathbf{G}_k + \gamma_k \left[ \partial_{\theta} L_c(\mathbf{Y}, \mathbf{X}^{(k)}, \boldsymbol{\theta}) - \mathbf{G}_k \right] \\ \mathbf{H}_{k+1} &= \mathbf{H}_k + \gamma_k \left[ \partial_{\theta}^2 L_c(\mathbf{Y}, \mathbf{X}^{(k)}, \boldsymbol{\theta}) \right. \\ &\quad \left. + \partial_{\theta} L_c(\mathbf{Y}, \mathbf{X}^{(k)}, \boldsymbol{\theta}) (\partial_{\theta} L_c(\mathbf{Y}, \mathbf{X}^{(k)}, \boldsymbol{\theta}))' - \mathbf{H}_k \right] \\ \mathbf{F}_{k+1} &= \mathbf{H}_{k+1} - \mathbf{G}_{k+1} (\mathbf{G}_{k+1})'. \end{aligned}$$

As the sequence  $(\hat{\boldsymbol{\theta}}^{(k)})_k$  converges to the maximum of an approximate likelihood, the sequence  $(\mathbf{F}_k)_k$  converges to the corresponding approximate Fisher information matrix. It is possible to initialize  $\mathbf{G}_0$  and  $\mathbf{H}_0$  to be a vector and a matrix of zeros respectively. We stress that we do not make use of the (approximate) Fisher information during the optimization.

## 4 Simulation studies

Simulations were coded in MATLAB (except for R examples using the `pomp` package) and executed on a Intel Core i7-2600 CPU 3.40 GhZ. Software code can be found online either at

<https://github.com/umbertopicchini/SAEM-ABC> or as supplementary material in the version of this paper published on *Computational Statistics* (doi:10.1007/s00180-017-0770-y), see also the footnote on page 2. For all examples we consider a Gaussian kernel for  $J_{j,\delta}$  as in (8). As described at the end of section 2.1, in SAEM we always set  $\gamma_k = 1$  for the first  $K_1$  iterations and  $\gamma_k = (k - K_1)^{-1}$  for  $k \geq K_1$  as in Lavielle [2014]. All results involving ABC are produced using algorithm 3 i.e. using trajectories selected via ABC-SMC. We compare our results with standard algorithms for Bayesian and “classical” inference, namely Gibbs sampling and particle marginal methods (PMM) [Andrieu and Roberts, 2009] for Bayesian inference and the improved iterated filtering (denoted in literature as IF2) found in Ionides et al. [2015] for maximum likelihood estimation. In order to perform a fair comparison between methods, we make use of well tested and maintained code to fit models with PMM and IF2 via the R `pomp` package [King et al., 2015]. All the methods mentioned above use sequential Monte Carlo algorithms (SMC), and their `pomp` implementation considers the bootstrap filter. We remark that our goal is not to consider specialized state-of-art SMC algorithms, with the notable exception mentioned below. Our focus is to compare the several inference methods above, while using the bootstrap filter for the trajectory proposal step: the bootstrap filter is also the approach considered in King et al. [2015] and Fasiolo et al. [2016] hence it is easier for us to compare methods using available software packages such as `pomp`. However in section 4.2 we also use a particles sampler that conditions upon data and is suitable for state-space models having latent process expressed by a stochastic differential equation [Golightly and Wilkinson, 2011].

## 4.1 Non-linear Gaussian state-space model

Consider the following Gaussian state-space model

$$\begin{cases} Y_j = X_j + \sigma_y \nu_j \\ X_j = 2 \sin(e^{X_{j-1}}) + \sigma_x \tau_j, \quad j = 1, \dots, n \end{cases} \quad (9)$$

with  $\nu_j, \tau_j \sim N(0, 1)$  i.i.d. and  $X_0 = 0$ . Parameters  $\sigma_x, \sigma_y > 0$  are the only unknowns and therefore we conduct inference for  $\theta = (\sigma_x^2, \sigma_y^2)$ .

We first construct the set of sufficient statistics corresponding to the complete log-likelihood  $L_c(\mathbf{Y}, \mathbf{X})$ . This is a very simple task since  $Y_j | X_j \sim N(X_j, \sigma_y^2)$  and  $X_j | X_{j-1} \sim N(2 \sin(e^{X_{j-1}}), \sigma_x^2)$  and therefore it is easy to show that  $S_{\sigma_x^2} = \sum_{j=1}^n (X_j - 2 \sin(e^{X_{j-1}}))^2$  and  $S_{\sigma_y^2} = \sum_{j=1}^n (Y_j - X_j)^2$  are sufficient for  $\sigma_x^2$  and  $\sigma_y^2$  respectively. By plugging these statistics into  $L_c(\mathbf{Y}, \mathbf{X})$  and equating to zero the gradient of  $L_c$  with respect to  $(\sigma_x^2, \sigma_y^2)$ , we find that the M-step of SAEM results in updated values for  $\sigma_x^2$  and  $\sigma_y^2$  given by  $S_{\sigma_x^2}/n$  and  $S_{\sigma_y^2}/n$  respectively. Expressions for the first, second and mixed derivatives, useful to obtain the Fisher information as in Section 3, are given in appendix.

### 4.1.1 Results

We generate  $n = 50$  observations for  $\{Y_j\}$  with  $\sigma_x^2 = \sigma_y^2 = 5$ , see Figure 1. We first describe results obtained using SAEM-ABC. Since the parameters of interest are positive, for numerical convenience we work on the log-transformed versions  $(\log \sigma_x, \log \sigma_y)$ . Our setup consists in running 30 independent experiments with SAEM-ABC: for each experiment we simulate parameter starting values for  $(\log \sigma_x, \log \sigma_y)$  independently generated from a bivariate Gaussian distribution with mean the true value of the parameters, i.e.  $(\log \sqrt{5}, \log \sqrt{5})$ , and diagonal covariance matrix having values (2,2) on its diagonal. For all 30 simulations we use the same data and the same setup except that in each simulation we use different starting values for the parameters. For each of the 30 experiments we let the threshold  $\delta$  decrease in the set of values  $\delta \in \{2, 1.7, 1.3, 1\}$  for a total of  $K = 400$  SAEM-ABC iterations, where we use  $\delta = 2$  for the first 80 iterations,  $\delta = 1.7$  for further 70 iterations,  $\delta = 1.3$  for further 50 iterations and  $\delta = 1$  for the remaining 200 iterations.




Figure 1: Nonlinear-Gaussian model: data when  $\sigma_x = \sigma_y = 2.23$ .

	$(M, \bar{M})=(500,200)$	$(M, \bar{M})=(1000,200)$	$(M, \bar{M})=(2000,200)$	$(M, \bar{M})=(1000,20)$
	$\sigma_x$ estimates (true = 2.23)			
SAEM-ABC	2.42 [2.39,2.47]	2.30 [2.27,2.32]	2.19 [2.17,2.24]	1.88 [1.81,1.93]
SAEM-SMC	2.54 [2.53,2.54]	2.55 [2.54,2.56]	2.55 [2.54,2.56]	1.99 [1.85,2.14]
IF2*	1.26 [1.21,1.41]	1.35 [1.28,1.41]	1.33 [1.28,1.40]	1.35 [1.28,1.41]
	$\sigma_y$ estimates (true = 2.23)			
SAEM-ABC	1.90 [1.87,1.94]	1.91 [1.88,1.95]	1.87 [1.84,1.93]	1.91 [1.89,1.96]
SAEM-SMC	0.11 [0.10,0.13]	0.06 [0.06,0.07]	0.04 [0.03,0.04]	1.23 [0.99,1.39]
IF2*	1.62 [1.56,1.75]	1.64 [1.58,1.67]	1.63 [1.59,1.67]	1.64 [1.58,1.67]

Table 1: Non-linear Gaussian model: medians and 1<sup>st</sup> – 3<sup>rd</sup> quartiles for estimates obtained on 30 independent simulations. (\*)The IF2 method resamples at every time point, while SAEM-ABC and SAEM-SMC resamples only when  $ESS < \bar{M}$ .

The influence of this choice is studied below. As explained in section 3, the largest value for  $\delta$  can be set intuitively, by looking at Figure 1, where it is apparent that considering deviations  $\delta = 2$  of the simulated observations  $Y_j^*$  from the actual observation  $Y_j$  should be reasonable. For example the empirical standard deviation of the differences  $|Y_j - Y_{j-1}|$  is about 2. Then we let  $\delta$  decrease progressively as SAEM-ABC evolves. We take  $K_1 = 300$  as the number of SAEM warmup iterations and use different numbers of particles  $M$  in our simulation studies, see Table 1. We impose resampling when the effective sample size ESS gets smaller than  $\bar{M}$ , for any attempted value of  $M$ . At first we show that taking  $\bar{M} = 200$  gives good results for SAEM-ABC but not for SAEM-SMC, see Table 1. Table 1 reports the median of the 30 estimates and their 1<sup>st</sup> – 3<sup>rd</sup> quartiles: we notice that  $M = 1,000$  particles are able to return satisfactory estimates when using SAEM-ABC. Figure 2 shows the rapid convergence of the algorithm towards the true parameter values for all the 30 repetitions (though difficult to notice visually, several simulations start at locations very far from the true parameter values). Notice that it only required about 150 seconds to perform all 30 simulations: this is the useful return out of the effort of constructing the analytic quantities necessary to run SAEM. Also, the algorithm is not very sensitive to the choice of  $\delta$ 's. For example, we also experimented with  $\delta \in \{4, 3, 2, 1\}$  and we obtained very similar results, for example our thirty experiments with  $(M, \bar{M}) = (1000, 200)$  resulted in medians (1st-3rd quartiles)  $\hat{\sigma}_x = 2.30$  [2.27,2.35],  $\hat{\sigma}_y = 1.90$  [1.86,1.91], which are very close to the ones in Table 1. Finally, notice that results are not overly sensitive to the way  $\delta$  is decreased: for example, if we let  $\delta$  decrease uniformly with steps of size  $1/3$ , that is  $\delta \in \{2, 1.67, 1.33, 1\}$  with  $\delta = 2$  for the first 50 iterations, then let it decrease every 50 iterations until  $\delta = 1$ , we obtain (when  $M = 1,000$  and  $\bar{M} = 200$ )  $\hat{\sigma}_x = 2.30$  [2.27,2.33],  $\hat{\sigma}_y = 1.91$  [1.87,1.95], compare with Table 1.

We then perform 30 simulations with SAEM-SMC using the same simulated data and parameters starting values as for SAEM-ABC. As from Table 1 simulations for  $\sigma_y$  converge to completely wrong values. For this case we also experimented with  $M = 5,000$  using  $\bar{M} = 2,000$




Figure 2: Non-linear Gaussian model: traces obtained with SAEM-ABC when using  $M = 1,000$  particles,  $\bar{M} = 200$  and  $(\sigma_x, \sigma_y) = (2.23, 2.23)$ . Horizontal lines are the true parameter values.

but this does not solve the problem with SAEM-SMC, even if we let the algorithm start at the true parameter values. We noted that when using  $M = 2,000$  particles with SAEM-ABC and  $\bar{M} = 200$  the algorithm resamples every fourth observation, and in a generic iteration we observed an ESS of about 100 at the last time point. Under the same setup SAEM-SMC resampled at each time point and resulted in an ESS of about 10 at the last time point (a study on the implications of frequent resampling is considered in section 4.1.3). Therefore we now perform a further simulation study to verify whether using a smaller  $\bar{M}$  (hence reducing the number of times resampling is performed) can improve the performance of SAEM-SMC. Indeed, using for example  $\bar{M} = 20$  gives better results for SAEM-SMC, see Table 1 (there we only use  $M = 1,000$  for comparison between methods). This is further investigated in section 4.1.3.

#### 4.1.2 Comparison with iterative filtering and a pseudo-marginal Bayesian algorithm

We compare the results above with the iterated filtering methodology for maximum likelihood estimation (IF2, [Ionides et al., 2015](#)), using the R package `pomp` [[King et al., 2015](#)]. We provide `pomp` with the same data and starting parameter values as considered in SAEM-ABC and SAEM-SMC. We do not provide a detailed description of IF2 here: it suffices to say that in IF2 particles are generated for both parameters  $\theta$  (e.g. via perturbations using random walks) and for the systems state (using the bootstrap filter). Moreover, same as with ABC methods, a “temperature” parameter (to use an analogy with the simulated annealing optimization method) is let decrease until the algorithm “freezes” around an approximated MLE. This temperature parameter, here denoted with  $\epsilon$ , is decreased in  $\epsilon \in \{0.9, 0.7, 0.4, 0.3, 0.2\}$  which seems appropriate as explained below, where the first value is used for the first 500 iterations of IF2, then each of the remaining values is used for 100 iterations, for a total of 900 iterations. Notice that the tested version of `pomp` (v. 1.4.1.1) uses a bootstrap filter that resamples at every time point, hence results obtained with IF2 are not directly comparable with SAEM-ABC and SAEM-SMC. Results are in Table 1 and a sample output from one of the simulations obtained with  $M = 1,000$  is in Figure 3. From the loglikelihood in Figure 3 we notice that the last major improvement in likelihood maximization takes place at iteration 600 when  $\epsilon$  becomes  $\epsilon = 0.7$ . Reducing  $\epsilon$  further does not give any additional benefit (we have verified this in a number of experiments with this model). Notice that in order to run, say, 400 iterations of IF2 with  $M = 1,000$  for a *single* experiment, instead of thirty, it required about 70 seconds. That is IF2 is about fourteen times slower than SAEM-ABC, although the comparison is not completely objective as we coded SAEM-ABC with Matlab, while IF2 is provided in an R package with forward model simulation implemented in C.

Finally we use a particle marginal method (PMM, [Andrieu and Roberts, 2009](#)) on a single

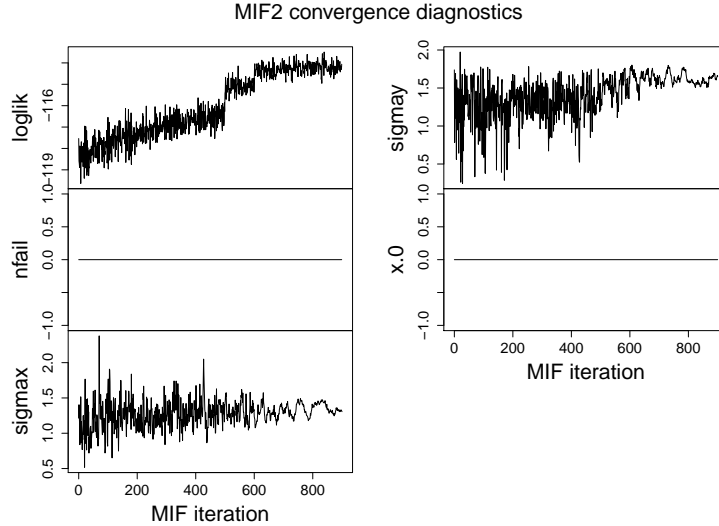


Figure 3: Non-linear Gaussian model: traces obtained with IF2 when using  $M = 1,000$  particles. (Top left) evolution of the loglikelihood function; (bottom left) evolution of  $\sigma_x$ ; (top right) evolution of  $\sigma_y$ .




Figure 4: Nonlinear-Gaussian model: ESS for SAEM-SMC when  $(\bar{M}, M) = (20, 1000)$  (solid line) and  $(\bar{M}, M) = (200, 1000)$  (dashed line) as a function of time.

simulation (instead of thirty), as this is a fully Bayesian methodology and results are not directly comparable with SAEM nor IF2. The PMM we construct approximates the likelihood function of the state space model using a bootstrap filter with  $M$  particles, then plugs such likelihood approximation into a Metropolis-Hastings algorithm. As remarkably shown in [Beaumont \[2003\]](#) and [Andrieu and Roberts \[2009\]](#), a PMM returns a Markov chain for  $\theta$  having the posterior  $\pi(\theta|\mathbf{Y}_{1:n})$  as its stationary distribution. This implies that PMM is an algorithm producing exact Bayesian inference for  $\theta$ , regardless the number of particles used to approximate the likelihood.

Once more we make use of facilities provided in `pomp` to run PMM. We set uniform priors  $U(0.1, 15)$  for both  $\sigma_x$  and  $\sigma_y$  and run 4,000 MCMC iterations of PMM using 2,000 particles for the likelihood approximation. Also, we set the PMM algorithm in the most favourable way, by starting it at the true parameter values (we are only interested in the inference results, rather than showing the performance of the algorithm when starting from remote values). The proposal function for the parameters uses an adaptive MCMC algorithm based on Gaussian random walks, and was tuned to achieve the optimal 7% acceptance rate [[Sherlock et al., 2015](#)]. For this single simulation, we obtained the following posterior means and 95% intervals:  $\hat{\sigma}_x = 1.52$  [0.42,2.56],  $\hat{\sigma}_y = 1.53$  [0.36,2.34].




Figure 5: Nonlinear-Gaussian model: number of distinct particles for SAEM-SMC when  $(\bar{M}, M) = (20, 1000)$  (solid line) and  $(\bar{M}, M) = (200, 1000)$  (dashed line) as a function of time.

### 4.1.3 The particles impoverishment problem

Figure 4 reports the effective sample size ESS as a function of time  $t$  (at a generic iteration of SAEM-SMC, the 20th in this case). As expected, for a smaller value of  $\bar{M}$  the ESS is most of times smaller than when a larger  $\bar{M}$  is chosen, with the exception of a few peaks. This is a direct consequence of performing resampling more frequently when  $\bar{M}$  is larger. However, a phenomenon that is known to be strictly linked to resampling is that of “samples impoverishment”, that is the resampling step reduces the “variety” of particles, by duplicating the ones with larger weight and killing the others. In fact, when many particles have a common “parent” at time  $t$ , these are likely to end close to each other at time  $t + 1$ . This has a negative impact on the inference because the purpose of the particles is to approximate the density (5) (or (6)) which generates the trajectory sampled at step 2 in algorithms 3 and 4. Lack of variety in the particles reduces the quality of this approximation.

Indeed, Figure 5 shows that the variety of the particles (as measured by the number of distinct particles) gets impoverished for a larger  $\bar{M}$ , notice for example that the solid line in Figure 5 almost always reaches its maximum attainable value  $M = 1,000$ , that is all particles are distinct, while this is often not the case for the dashed line. Since the trajectory  $\mathbf{X}^{(k)}$  that is selected at iteration  $k$  of SAEM, either in step 2 of algorithm 3 or in step 2 of algorithm 4, follows from backward-tracing the genealogy of a certain particle, having variety in the cloud of particles is crucial here.

This seems related to the counter-intuitive good performance of the ABC-filter, even though SAEM-ABC is based on the additional approximation induced by the  $J$  function in equation (8). We now produce plots for the ESS values and the number of distinct particles at the smallest value of the ABC threshold  $\delta$ . As we see in Figure 6, while the ESS are not much different from the ones in Figure 4, instead the number of distinct particles in Figure 7 is definitely higher than the SAEM-SMC counterpart in Figure 5, meaning that such number drops below the maximum  $M = 1,000$  fewer times. For example from Figure 5 we can see that when  $(\bar{M}, M) = (200, 1000)$  the number of distinct particles drops 19 times away from the maximum  $M = 1,000$ , when using SAEM-SMC. With SAEM-ABC this number drops only 13 times (Figure 7). When the number of resampling steps is reduced, using  $(\bar{M}, M) = (20, 1000)$ , we have more even results, with the number of distinct particles dropping six times for SAEM-SMC and five times for SAEM-ABC.

As a support to this remark, refer to Table 2: we perform thirty independent estimation procedures, and in each we obtain the sample mean of the ESS (means computed over varying time  $t$ ) and the sample mean of the number of distinct particles (again over varying  $t$ ). Then, we report the mean over the thirty estimated sample means of the ESS (and corresponding standard deviation) and the same for the number of distinct particles. Clearly numbers are favourable to SAEM-ABC, showing consistently larger values (with small variation between the thirty experiments).

We argue that when many particles, as generated in the bootstrap filter, fail to get close






Figure 6: Nonlinear-Gaussian model: ESS for SAEM-ABC when  $(\bar{M}, M) = (20, 1000)$  (solid line) and  $(\bar{M}, M) = (200, 1000)$  (dashed line) as a function of time.




Figure 7: Nonlinear-Gaussian model: number of distinct particles for SAEM-ABC when  $(\bar{M}, M) = (20, 1000)$  (solid line) and  $(\bar{M}, M) = (200, 1000)$  (dashed line) as a function of time.

to the target observations, then resampling is frequently triggered, this degrading the variety of the particles. However with an ABC filter, particles receive some additional weighting (due to a  $J_{j,\delta} > 1$  in (8) when  $Y_j^{*(m)} \approx Y_j$  and a small  $\delta$ ) which allows for a larger number of particles to have a non-negligible weight. While it is *not true* that an ABC filter is in general expected to perform better than a non-ABC one, here we find that the naive bootstrap filter performs worse than the ABC counterpart for the reasons discussed above.

#### 4.1.4 Relation with Gibbs sampling

As previously mentioned, the generation of  $\mathbf{s}_k$  in (3) followed by corresponding parameter estimates  $\hat{\boldsymbol{\theta}}^{(k)}$  in (4) is akin to two steps of a Gibbs sampling algorithm, with the important distinction that in SAEM  $\hat{\boldsymbol{\theta}}^{(k)}$  is produced by a deterministic step (for given  $\mathbf{s}_k$ ). We show that the construction of a Gibbs-within-Metropolis sampler is possible, but that a naive implementation fails while a “non-central parametrization” seems necessary [Papaspiliopoulos et al., 2007]. For given initial vectors  $(\mathbf{X}^{(0)}, \sigma_x^{(0)}, \sigma_y^{(0)})$ , we alternate sampling from the conditional distributions (i)

	$(M, \bar{M}) = (1000, 20)$		$(M, \bar{M}) = (1000, 200)$	
	ESS	# distinct particles	ESS	# distinct particles
SAEM-SMC	202.20 (61.70)	871.83 (90.62)	252.10 (83.31)	616.40 (212.61)
SAEM-ABC	202.62 (12.31)	905.64 (7.95)	351.80 (14.75)	812.85 (4.68)

Table 2: Nonlinear-Gaussian model: mean ESS and mean number of distinct particles (and corresponding standard deviations in parentheses) at a generic iteration of SAEM-ABC and SAEM-SMC. Averages and standard deviations are taken over 30 independent repetitions of the experiment.




Figure 8: Nonlinear Gaussian model: three chains with different starting values from the Metropolis-within-Gibbs sampler for data generated with  $(\sigma_x, \sigma_y) = (2.23, 2.23)$ . Horizontal lines are the true parameter values.

$p(\mathbf{X}^{(b)} | \sigma_x^{(b-1)}, \sigma_y^{(b-1)}, \mathbf{Y})$ , (ii)  $p(\sigma_x^{(b)} | \sigma_y^{(b-1)}, \mathbf{X}^{(b)}, \mathbf{Y})$  and (iii)  $p(\sigma_y^{(b)} | \sigma_x^{(b)}, \mathbf{X}^{(b)}, \mathbf{Y})$  where  $b$  represents the iteration counter in the Gibbs sampler ( $b = 1, 2, \dots$ ). The resulting multivariate sample  $(\mathbf{X}^{(b)}, \sigma_x^{(b)}, \sigma_y^{(b)})$  is a draw from the posterior distribution  $\pi(\mathbf{X}, \sigma_x, \sigma_y | \mathbf{Y})$ . We cannot sample from the conditional densities in (i)–(iii), however at a generic iteration  $b$  it is possible to incorporate a single Metropolis-Hastings step targeting the corresponding densities in (i)–(iii) separately, resulting in a Metropolis-within-Gibbs sampler, see e.g. Liu [2008]. Notice that what (i) implies is a joint sampling (block-update) for all the elements in  $\mathbf{X}$ , however it is also possible to sample individual elements one at a time (single-site update) from  $p(X_j^{(b)} | X_{j-1}^{(b)}, X_{j+1}^{(b-1)}, \sigma_x^{(b-1)}, \sigma_y^{(b-1)}, \mathbf{Y})$ . For both single-site and block-update sampling, the mixing of the resulting chain is very poor, this resulting from the high correlations between sampled quantities, notably the correlation between elements in  $\mathbf{X}$  and also between  $\mathbf{X}$  and  $(\sigma_x, \sigma_y)$ . However, there exists a simple solution based on breaking down the dependence between some of the involved quantities: for example at iteration  $b$  we propose in block a vector  $\mathbf{X}^\#$  generated “blindly” (i.e. conditionally on the current parameter values  $(\sigma_x^{(b-1)}, \sigma_y^{(b-1)})$ , but unconditionally on  $\mathbf{Y}$ ) by iterating through (9) then accept/reject the proposal according to a Metropolis-Hastings step, then define  $\mathbf{X}^{*(b)} := \tilde{\mathbf{X}}/\sigma_x^{(b-1)}$ , where  $\tilde{\mathbf{X}}$  is the last accepted proposal of  $\mathbf{X}$  (that is  $\tilde{\mathbf{X}} \equiv \mathbf{X}^\#$  if  $\mathbf{X}^\#$  has been accepted). Sample from (ii)  $p(\sigma_x^{(b)} | \sigma_y^{(b-1)}, \mathbf{X}^{*(b)}, \mathbf{Y})$  and (iii)  $p(\sigma_y^{(b)} | \sigma_x^{(b)}, \mathbf{X}^{*(b)}, \mathbf{Y})$ , and transform back to  $\tilde{\mathbf{X}} := \sigma_x^{(b)} \mathbf{X}^{*(b)}$ . This type of updating scheme is known as non-central parametrization [Papaspiliopoulos et al., 2007]. The expressions for the conditional densities in (i)–(iii) are given in the appendix for the interested reader. Note that the SAEM-ABC and SAEM-SMC do not require a non-central parametrization. The maximization step of SAEM smooths out the correlation between the proposed parameter and the latent state, and the numerical convergence of the algorithms still occur (this has also been noticed in previous papers on SAEM-MCMC, see Kuhn and Lavielle, 2005, Donnet and Samson, 2008).

Once more, we attempt estimating data produced with  $(\sigma_x, \sigma_y) = (2.23, 2.23)$  (as in section 4.1.1) this time using Metropolis-within-Gibbs. Trace plots in Figures 8 show satisfactory mixing for three chains starting at three random values around  $(\sigma_x, \sigma_y) = (6, 6)$ . However, while the chains initialized at values far from the truth rapidly approach the true values, the 95% posterior intervals fail to include them, see Figure 9. If we re-execute the same experiment, this time with data generated with smaller noise  $(\sigma_x, \sigma_y) = (0.71, 0.71)$ , it seems impossible to catch the ground truth parameters. From a typical chain, we obtain the following posterior means and 95% posterior intervals:  $\hat{\sigma}_x = 1.59$  [1.18, 2.02],  $\hat{\sigma}_y = 1.63$  [1.28, 2.08]. Clearly for noisy time-dependent data, such as state-space models, particles based methods seem to better address cases where noisy data are affected by both measurement and systemic noise.

## 4.2 A pharmacokinetics model

Here we consider a model for pharmacokinetics dynamics. For example we could formulate a model to study the Theophylline drug pharmacokinetics. This example has often been described




Figure 9: Nonlinear Gaussian model: marginal posteriors from the Metropolis-within-Gibbs sampler for data generated with  $(\sigma_x, \sigma_y) = (2.23, 2.23)$ . Vertical lines are the true parameter values.

in literature devoted to longitudinal data modelling with random parameters (mixed-effects models), see [Pinheiro and Bates \[1995\]](#) and [Donnet and Samson \[2008\]](#). Same as in [Picchini \[2014\]](#) here we do not consider a mixed-effects model. We denote with  $X_t$  the level of drug concentration in blood at time  $t$  (hrs). Consider the following non-autonomous stochastic differential equation:

$$dX_t = \left( \frac{Dose \cdot K_a \cdot K_e}{Cl} e^{-K_a t} - K_e X_t \right) dt + \sigma \sqrt{X_t} dW_t, \quad t \geq t_0 \quad (10)$$

where  $Dose$  is the known drug oral dose received by a subject,  $K_e$  is the elimination rate constant,  $K_a$  the absorption rate constant,  $Cl$  the clearance of the drug and  $\sigma$  the intensity of the intrinsic stochastic noise. We simulate data at  $n = 100$  equispaced sampling times  $\{t_1, t_\Delta, \dots, t_{100\Delta}\} = \{1, 2, \dots, 100\}$  where  $\Delta = t_j - t_{j-1} = 1$ . The drug oral dose is chosen to be 4 mg. After the drug is administered, we consider as  $t_0 = 0$  the time when the concentration first reaches  $X_{t_0} = X_0 = 8$ . The error model is assumed to be linear,  $Y_j = X_j + \varepsilon_j$  where the  $\varepsilon_j \sim N(0, \sigma_\varepsilon^2)$  are i.i.d.,  $j = 1, \dots, n$ . Inference is based on data  $\{Y_1, \dots, Y_n\}$  collected at corresponding sampling times. Parameter  $K_a$  is assumed known, hence parameters of interest are  $\boldsymbol{\theta} = (K_e, Cl, \sigma^2, \sigma_\varepsilon^2)$  as  $X_0$  is also assumed known.

Equation (10) has no available closed-form solution, hence simulated data are created in the following way. We first simulate numerically a solution to (10) using the Euler-Maruyama discretization with stepsize  $h = 0.05$  on the time interval  $[t_0, 100]$ . The Euler-Maruyama scheme is given by

$$X_{t+h} = X_t + \left( \frac{Dose \cdot K_a \cdot K_e}{Cl} e^{-K_a t} - K_e X_t \right) h + (\sigma \sqrt{h \cdot X_t}) Z_{t+h}$$

where the  $\{Z_t\}$  are i.i.d.  $N(0, 1)$  distributed. The grid of generated values  $\mathbf{X}_{0:N}$  is then linearly interpolated at sampling times  $\{t_1, \dots, t_{100}\}$  to give  $\mathbf{X}_{1:n}$ . Finally residual error is added to  $\mathbf{X}_{1:n}$  according to the model  $Y_j = X_j + \varepsilon_j$  as explained above. Since the errors  $\varepsilon_j$  are independent, data  $\{Y_j\}$  are conditionally independent given the latent process  $\{X_t\}$ .

**Sufficient statistics for SAEM** The complete likelihood is given by

$$p(\mathbf{Y}, \mathbf{X}_{0:N}; \boldsymbol{\theta}) = p(\mathbf{Y} | \mathbf{X}_{0:N}; \boldsymbol{\theta}) p(\mathbf{X}_{0:N}; \boldsymbol{\theta}) = \prod_{j=1}^n p(Y_j | X_j; \boldsymbol{\theta}) \prod_{i=1}^N p(X_i | X_{i-1}; \boldsymbol{\theta})$$

where the unconditional density  $p(x_0)$  is disregarded in the last product since we assume  $X_0$  deterministic. Hence the complete-data loglikelihood is

$$L_c(\mathbf{Y}, \mathbf{X}_{0:N}; \boldsymbol{\theta}) = \sum_{j=1}^n \log p(Y_j | X_j; \boldsymbol{\theta}) + \sum_{i=1}^N \log p(X_i | X_{i-1}; \boldsymbol{\theta}).$$

Here  $p(y_j|x_j; \boldsymbol{\theta})$  is a Gaussian with mean  $x_j$  and variance  $\sigma_\varepsilon^2$ . The transition density  $p(x_i|x_{i-1}; \boldsymbol{\theta})$  is not known for this problem, hence we approximate it with the Gaussian density induced by the Euler-Maruyama scheme, that is

$$p(x_i|x_{i-1}; \boldsymbol{\theta}) \approx \frac{1}{\sigma\sqrt{2\pi x_{i-1}h}} \exp\left\{-\frac{[x_i - x_{i-1} - (\frac{Dose \cdot K_a \cdot K_e}{Cl} e^{-K_a \tau_{i-1}} - K_e x_{i-1})h]^2}{2\sigma^2 x_{i-1}h}\right\}. \quad (11)$$

Notice the Gaussian distribution implied by (11) shares some connection with tools developed for optimal states predictions in signal processing, such as the unscented Kalman filter, e.g. [Sitz et al. \[2002\]](#). We now derive sufficient summary statistics for the parameters of interest, based on the complete loglikelihood. Regarding  $\sigma_\varepsilon^2$  this is trivial as we only have to consider  $\sum_{j=1}^n \log p(y_j|x_j; \boldsymbol{\theta})$  to find that a sufficient statistic is  $S_{\sigma_\varepsilon^2} = \sum_{j=1}^n (y_j - x_j)^2$ . Regarding the remaining parameters we have to consider  $\sum_{i=1}^N \log p(x_i|x_{i-1}; \boldsymbol{\theta})$ . For  $\sigma^2$  a sufficient statistic is

$$S_{\sigma^2} = \sum_{i=1}^N \left( \frac{[x_i - x_{i-1} - (\frac{Dose \cdot K_a \cdot K_e}{Cl} e^{-K_a \tau_{i-1}} - K_e x_{i-1})h]^2}{x_{i-1}h} \right).$$

Regarding  $K_e$  and  $Cl$  reasoning is a bit more involved: we can write

$$\begin{aligned} \sum_{i=1}^N \log p(x_i|x_{i-1}; \boldsymbol{\theta}) &\propto \sum_{i=1}^N \frac{[x_i - x_{i-1} - (\frac{Dose \cdot K_a \cdot K_e}{Cl} e^{-K_a \tau_{i-1}} - K_e x_{i-1})h]^2}{x_{i-1}} \\ &= \sum_{i=1}^N \left[ \frac{x_i - x_{i-1}}{\sqrt{x_{i-1}}} - \left( \frac{Dose \cdot K_a \cdot K_e}{Cl \sqrt{x_{i-1}}} e^{-K_a \tau_{i-1}} - \frac{K_e x_{i-1}}{\sqrt{x_{i-1}}} \right) h \right]^2. \end{aligned}$$

The last equality suggests a linear regression approach  $E(V) = \beta_1 C_1 + \beta_2 C_2$  for ‘‘responses’’  $V_i = (x_i - x_{i-1})/\sqrt{x_{i-1}}$  and ‘‘covariates’’

$$\begin{aligned} C_{i1} &= \frac{Dose \cdot K_a e^{-K_a \tau_{i-1}} h}{\sqrt{x_{i-1}}} \\ C_{i2} &= -\frac{x_{i-1}}{\sqrt{x_{i-1}}} h = -\sqrt{x_{i-1}} h \end{aligned}$$

and  $\beta_1 = K_e/Cl$ ,  $\beta_2 = K_e$ . By considering the design matrix  $\mathbf{C}$  with columns  $\mathbf{C}_1$  and  $\mathbf{C}_2$ , that is  $\mathbf{C} = [\mathbf{C}_1, \mathbf{C}_2]$ , from standard regression theory we have that  $\hat{\boldsymbol{\beta}} = (\mathbf{C}'\mathbf{C})^{-1}\mathbf{C}'\mathbf{V}$  is a sufficient statistic for  $\boldsymbol{\beta} = (\beta_1, \beta_2)$ , where  $'$  denotes transposition. We take  $S_{K_e} := \hat{\beta}_2$  also to be used as the updated value of  $K_e$  in the maximisations step of SAEM. Then we have that  $\hat{\beta}_1$  is sufficient for the ratio  $K_e/Cl$  and use  $\hat{\beta}_2/\hat{\beta}_1$  as the update of  $Cl$  in the M-step of SAEM. The updated values of  $\sigma$  and  $\sigma_\varepsilon$  are given by  $\sqrt{S_{\sigma^2}/N}$  and  $\sqrt{S_{\sigma_\varepsilon^2}/n}$  respectively.

### 4.3 Results

We consider an experiment where 50 datasets of length  $n = 100$  each are independently simulated using parameter values  $(K_e, K_a, Cl, \sigma, \sigma_\varepsilon) = (0.05, 1.492, 0.04, 0.1, 0.1)$ . All results pertaining SAEM-ABC use the Gaussian kernel (8). For SAEM-ABC we use a number of ‘‘schedules’’ to decrease the threshold  $\delta$ . One of our attempts decreases the threshold as  $\delta \in \{0.5, 0.2, 0.1, 0.03\}$  (results for this schedule are reported as SAEM-ABC(1) in Table 3): same as in the previous application, all we need is to determine an appropriate order of magnitude for the largest  $\delta$ . As an heuristic, we have that the empirical standard deviation of the differences  $|Y_j - Y_{j-1}|$  is about 0.3. The first value of  $\delta$  is used for the first 80 iterations then it is progressively decreased every 50 iterations. For both SAEM-ABC and SAEM-SMC we use  $K_1 = 250$  and  $K = 300$  and optimization started at parameter values very far from the true values: starting

values are  $K_e = 0.80$ ,  $Cl = 10$ ,  $\sigma = 0.14$  and  $\sigma_\varepsilon = 1$ . We first show results using  $(M, \bar{M}) = (200, 10)$ . We start with SAEM-ABC, see Figure 10 where the effect of using a decreasing  $\delta$  is evident, especially on the trajectories for  $\sigma_\varepsilon$ . All trajectories but a single erratic one converge towards the true parameter values. Estimation results are in Table 3, giving results for three different decreasing schedules for  $\delta$ , reported as SAEM-ABC(0) for  $\delta \in \{0.5, 0.2, 0.1, 0.05, 0.01\}$ , the already mentioned SAEM-ABC(1) having  $\delta \in \{0.5, 0.2, 0.1, 0.03\}$  and finally SAEM-ABC(2) where  $\delta \in \{1, 0.4, 0.1\}$ . Results are overall satisfactory for all parameters but  $\sigma$ , which remains unidentified for all attempted SAEM algorithms, including those discussed later. The benefit of decreasing  $\delta$  to values close to zero are noticeable, that is among the algorithms using ABC, SAEM-ABC(0) gives the best results.

From results in Table 3 and Figure 11, again considering the case  $M = 200$ , we notice that SAEM-SMC struggles in identifying most parameters with good precision. As discussed later on, when showing results with  $M = 30$ , it is clear that when a very limited amount of particles is available (which is of interest when it is computationally demanding to forward-simulate from a complex model) SAEM-SMC is suboptimal compared to SAEM-ABC at least for the considered example. However results improve noticeably for SAEM-SMC as soon as the number of particles is enlarged, say to  $M = 1,000$ . In fact for  $M = 1,000$  the inference results for SAEM-ABC(0) and SAEM-SMC are basically the same. It is of course interesting to uncover the reason why SAEM-ABC(0) performs better than SEM-SMC when the number of particles is small, e.g.  $M = 200$  (see later on for results with  $M = 30$ ). In Figure 12 we compare approximations to the distribution  $\pi(\mathbf{X}_t | \mathbf{Y}_{1:t-1})$ , that is the distribution of the state at time  $t$  given previous data, and the filtering distribution  $\pi(\mathbf{X}_t | \mathbf{Y}_{1:t})$  including the most recent data. The former one,  $\pi(\mathbf{X}_t | \mathbf{Y}_{1:t-1})$ , is approximated by kernel smoothing applied on the particles  $\mathbf{X}_t^{(m)}$ . The filtering distribution  $\pi(\mathbf{X}_t | \mathbf{Y}_{1:t})$  is approximated by kernel smoothing assigning weights  $W_t^{(m)}$  to the corresponding particles. It is from the (particles induced) discrete distribution approximating  $\pi(\mathbf{X}_t | \mathbf{Y}_{1:t})$  that particles are sampled when propagating to the next time point. Both densities are computed from particles obtained at the last SAEM iteration,  $K$ , for different values of  $t$ , at the beginning of the observational interval ( $t = 10$ ), at  $t = 40$  and towards the end of the observational interval ( $t = 70$ ). Since in Figure 12 we also report the true values of  $\mathbf{X}_t$ , we can clearly see that, for SAEM-SMC, as  $t$  increases the true value of  $\mathbf{X}_t$  is unlikely under  $\pi(\mathbf{X}_t | \mathbf{Y}_{1:t})$  (the orange curve). In particular, by looking at the orange curve in panel (f) in Figure 12 we notice that many of the particles ending-up far from the true  $\mathbf{X}_t = 1.8$  receive non-negligible weight. Instead in panel (e) only particles very close to the true value of the state receive considerable weight (notice the different scales on the abscissas for panel (e) and (f)). Therefore it is not unlikely that for SAEM-SMC several ‘‘remote’’ particles are resampled and propagated. The ones resampled in SAEM-ABC received a weight  $W_t^{(m)} \propto J_{t,\delta}$  which is large only if the simulated observation is very close to the actual observation, since  $\delta$  is very small. Therefore for SAEM-SMC, when  $M$  is small, the path sampled in step 2 of algorithm 4 (which is from  $\pi(\mathbf{X}_{1:t_n} | \mathbf{Y}_{1:t_n})$ ) is poor and the resulting inference biased. For a larger  $M$  (e.g.  $M = 1,000$ ) SAEM-SMC enjoys a larger number of opportunities for particles to fall close to the targeted observation hence an improved inference.

An improvement over the bootstrap filter used in SAEM-SMC is given by a methodology where particles are not proposed from the transition density of the latent process (the latter proposes particles ‘‘blindly’’ with respect to the next data point). For example, [Golightly and Wilkinson \[2011\]](#) consider a proposal distribution based on a diffusion bridge, conditionally to observed data. Their methodology is specific for state-space models driven by a stochastic differential equation whose approximate solution is obtained via the Euler-Maruyama discretisation, which is what we require. We use their approach to ‘‘propagate forward’’ particles. We write SAEM-GW to denote a SAEM algorithm using the proposal sampler in [Golightly and Wilkinson \[2011\]](#). By looking at Table 3 we notice the improvement over the simpler SAEM-SMC when  $M = 200$ . In fact SAEM-GW gives the best results of all SAEM-based algorithms we attempted,




Figure 10: Theophylline model: 50 independent estimations using  $K = 300$  iterations of SAEM-ABC when  $M = 200$  and  $\delta \in \{0.5, 0.2, 0.1, 0.03\}$ . Top:  $K_e$  (left) and  $Cl$  (right). Bottom:  $\sigma$  (left) and  $\sigma_\epsilon$  (right). Horizontal lines are the true parameter values.

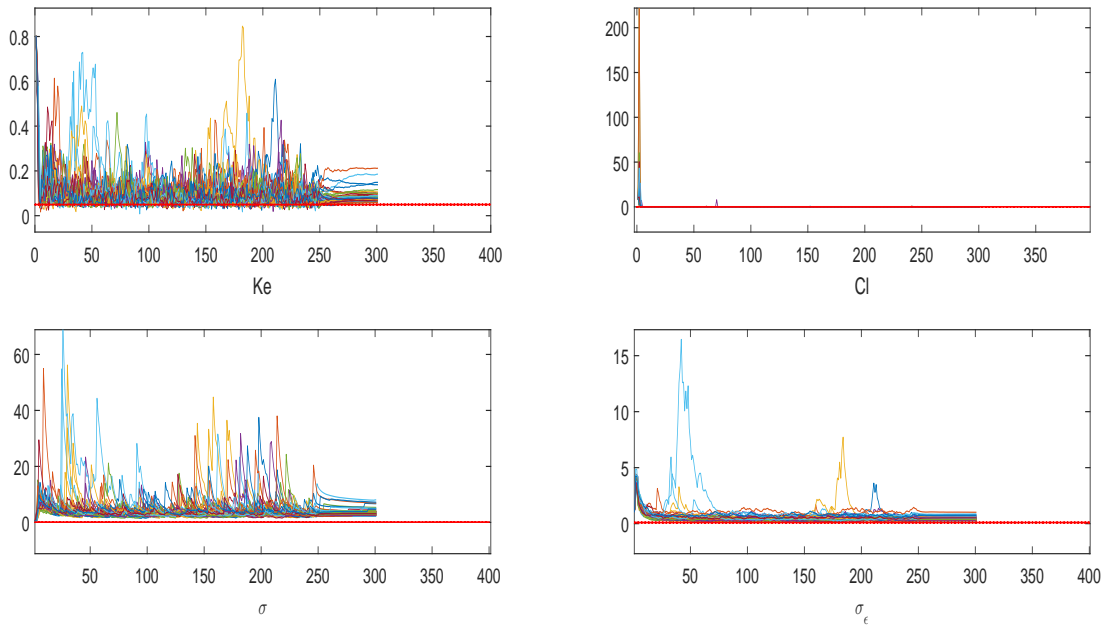


Figure 11: Theophylline model: 50 independent estimations using  $K = 300$  iterations of SAEM-SMC when  $M = 200$ . Top:  $K_e$  (left) and  $Cl$  (right). Bottom:  $\sigma$  (left) and  $\sigma_\epsilon$  (right). Horizontal lines are the true parameter values.




Figure 12: Kernel smoothed approximations of  $\pi(\mathbf{X}_t | \mathbf{Y}_{1:t-1})$  (blue) and  $\pi(\mathbf{X}_t | \mathbf{Y}_{1:t})$  (orange) at different values of  $t$  for the case  $(M = 200, \bar{M} = 10)$ . Green asterisks are the true values of  $\mathbf{X}_t$ . Notice scales on x-axes are different.

$(M, \bar{M}) = (200, 10)$							
	$K_e$		$Cl$		$\sigma$		$\sigma_\varepsilon$
true values	0.050		0.040		0.1		0.1
SAEM-ABC(0)	0.059 [0.054,0.067]		0.034 [0.027,0.038]		1.57 [1.35,2.15]		0.15 [0.11,0.22]
SAEM-ABC(1)	0.063 [0.058,0.074]		0.031 [0.026,0.035]		2.09 [1.56,2.80]		0.16 [0.13,0.25]
SAEM-ABC(2)	0.073 [0.065,0.088]		0.025 [0.022,0.028]		2.69 [2.44,3.50]		0.32 [0.27,0.37]
SAEM-SMC	0.078 [0.071,0.087]		0.022 [0.019,0.027]		3.45 [3.01,4.09]		0.45 [0.36,0.55]
SAEM-GW	0.061 [0.057,0.066]		0.032 [0.026,0.036]		1.84 [1.44,2.81]		0.12 [0.09,0.21]

$(M, \bar{M}) = (1000, 100)$							
	$K_e$		$Cl$		$\sigma$		$\sigma_\varepsilon$
true values	0.050		0.040		0.1		0.1
SAEM-ABC(0)	0.061 [0.054,0.065]		0.033 [0.027,0.035]		1.78 [1.48,2.48]		0.13 [0.09,0.19]
SAEM-ABC(1)	0.063 [0.059,0.070]		0.031 [0.025,0.035]		1.93 [1.65,2.41]		0.15 [0.13,0.23]
SAEM-ABC(2)	0.073 [0.064,0.081]		0.024 [0.021,0.028]		2.98 [2.54,3.42]		0.27 [0.24,0.35]
SAEM-SMC	0.062 [0.057,0.068]		0.033 [0.027,0.036]		1.91 [1.40,2.25]		0.12 [0.09,0.19]
SAEM-GW	0.061 [0.057,0.068]		0.032 [0.025,0.035]		1.90 [1.40,2.68]		0.13 [0.1,0.24]

Table 3: Theophylline: medians and 1<sup>st</sup> – 3<sup>rd</sup> quartiles for estimates obtained on 50 independent simulations using SAEM-ABC and SAEM-SMC. SAEM-ABC(0) denotes results obtained with  $\delta \in \{0.5, 0.2, 0.1, 0.05, 0.01\}$ , SAEM-ABC(1) denotes results obtained with  $\delta \in \{0.5, 0.2, 0.1, 0.03\}$  and SAEM-ABC(2) denotes results obtained with  $\delta \in \{1, 0.4, 0.1\}$ .

the downside being that such sampler is not very general and is only applicable to a specific class of models, namely (i) state-space models having an SDE that has to be numerically solved via Euler-Maruyama, (ii) observations having additive Gaussian noise, and (iii) observations having a state entering linearly, e.g.  $y = a \cdot x + \varepsilon$  for some constant  $a$ .

Same as in section 4.1, we consider Bayesian estimation using a particle marginal method (PMM). PMM is run with 1000 particles for 2,000 MCMC iterations. However it turns out that PMM cannot be initialized at the same remote starting values we used for SAEM, as the approximated log-likelihood function at the starting parameter results not-finite for the considered number of particles. Therefore we let PMM start at  $K_e = 0.05$ ,  $Cl = 0.04$ ,  $\sigma = 0.2$ ,  $\sigma_\varepsilon = 0.3$  and use priors  $K_e \sim U(0.01, 0.2)$ ,  $Cl \sim U(0.01, 0.2)$ ,  $\sigma \sim U(0.01, 0.3)$ ,  $\sigma_\varepsilon \sim U(0.05, 0.5)$ . Posterior means and 95% intervals are:  $K_e = 0.076$  [0.067,0.084],  $Cl = 0.056$  [0.047,0.068],  $\sigma = 0.13$  [0.10,0.16],  $\sigma_\varepsilon = 0.12$  [0.099,0.139]. Hence, for starting parameter values close to the true values PMM behaves well (and estimates  $\sigma$  correctly), but otherwise it might be impossible to initialize PMM, as also shown in Fasiolo et al. [2016].

We now explore whether the ABC approach can be of aid when saving computational time is essential, for example when simulating from the model is expensive. Although the model here considered can be simulated relatively quickly (it requires numerical integration hence computing times are affected by the size of the integration stepsize  $h$ ), assuming this is not the case we explore what could happen if we can only afford running a simulation with  $M = 30$  particles. We run 100 simulations with this setting. To ease graphic representation in the presence of outliers, estimates are reported on log-scales in the boxplots in Figures 13–15. We notice that, with such a small number of particles, SAEM-ABC is still able to estimate  $K_e$  and  $Cl$  accurately whereas SAEM-SMC returns more biased estimates for all parameters. Also, while both methodologies fail in estimating  $\sigma$  and  $\sigma_\varepsilon$  when  $M = 30$  SAEM-ABC is still better than SAEM-SMC confirming the previous finding that, should the computational budget be very limited, SAEM-ABC is a viable option. SAEM-GW is instead able to estimate also the residual error variability  $\sigma_\varepsilon$ .






Figure 13: Theophylline model: estimates obtained with SAEM-SMC when  $M = 30$ . From left to right:  $\log K_e$ ,  $\log Cl$ ,  $\log \sigma$  and  $\log \sigma_\epsilon$ . Green lines are the true parameter values on log-scale.

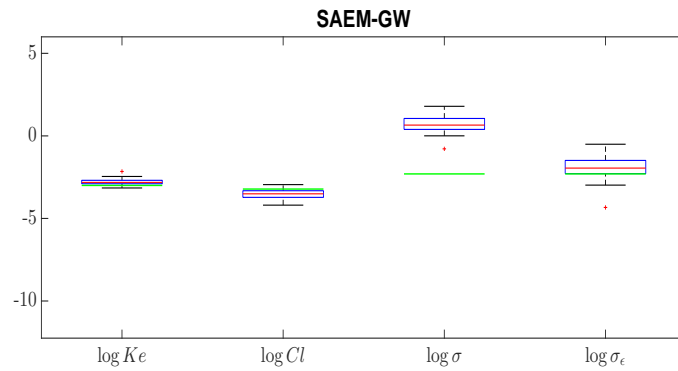


Figure 14: Theophylline model: estimates obtained with the Golightly-Wilkinson sampler coupled with SAEM (SAEM-GW) when  $M = 30$ . From left to right:  $\log K_e$ ,  $\log Cl$ ,  $\log \sigma$  and  $\log \sigma_\epsilon$ . Green lines are the true parameter values on log-scale.




Figure 15: Theophylline model: estimates obtained with SAEM-ABC when  $M = 30$ . From left to right:  $\log K_e$ ,  $\log Cl$ ,  $\log \sigma$  and  $\log \sigma_\epsilon$ . Green lines are the true parameter values on log-scale.

## 5 Summary

We have introduced a methodology for approximate maximum likelihood estimation of the parameters in state-space models, incorporating an approximate Bayesian computation (ABC) strategy. The general framework is the stochastic approximation EM algorithm (SAEM) of [De-lyon et al. \[1999\]](#), and we embed a sequential Monte Carlo ABC filter into SAEM. SAEM requires model-specific analytic computations, at the very least the derivation of sufficient statistics for the complete log-likelihood, to approach the parameters maximum likelihood estimate with minimal computational effort. However at any iteration of SAEM, it is required the availability of a filtered trajectory of the latent systems state, which we provide via the ABC filter of [Jasra et al. \[2012\]](#). We call this algorithm SAEM-ABC. An advantage of using the ABC filter is its flexibility, as it is possible to modify its setup to influence the weighting of the particles. In other words, it is possible to tune a positive tolerance  $\delta$  to enhance the importance of those particles that are the closest to the observations. This produced sampled trajectories for step 2 of algorithm 3 that resulted in a less biased parameter inference. This observation turned especially true for experiments run with a limited number of particles ( $M = 30$  or  $200$ ), which is relevant for computationally intensive models not allowing for the propagation of a large number of particles. We compared our SAEM-ABC algorithm with a version of SAEM employing the bootstrap filter of [Gordon et al. \[1993\]](#). The bootstrap filter is the simplest sequential Monte Carlo algorithm, and is typically a default option in many software packages (e.g. the `smfsb` and `pomp` R packages, [Wilkinson, 2015](#) and [King et al., 2015](#) respectively). In our work we show that, in some cases, SAEM-SMC (that is SAEM using a bootstrap filter) is sometimes inferior to SAEM-ABC, for example when the number of particles is small (section 4.3) or when too frequent resampling causes “particles impoverishment” (section 4.1.3). SAEM-ABC requires the user to specify a sequence of ABC thresholds  $\delta_1 > \dots > \delta_L > 0$ . For one-dimensional time-series setting these thresholds is intuitive, since these represent standard deviations of perturbed (simulated) observations. Therefore the size of the largest one ( $\delta_1$ ) can be determined by looking at plots of the observed time-series.

Ultimately, while we are not claiming that an ABC filter should in general be preferred to a non-ABC filter, as the former one induces some approximation, it can be employed when it is difficult to construct (or implement) a more advanced sequential Monte Carlo filter. In our second application we consider a sequential Monte Carlo filter due to [Golightly and Wilkinson \[2011\]](#). This filter (which we call GW) is specifically designed for state-space models driven by stochastic differential equations (SDEs) requiring numerical discretization. While GW improves noticeably over the basic bootstrap filter, GW is not very general: again, it is specific for state-space models driven by SDEs; the observation equation must have latent states entering linearly and measurement errors must be Gaussian distributed. The ABC approach instead does not impose any limitation on the model structure.

## Acknowledgements

We thank anonymous reviewers for helping improve considerably the quality of the work. Umberto Picchini was partially funded by the Swedish Research Council (VR grant 2013-5167). Adeline Samson was partially supported by the LabEx PERSYVAL-Lab (ANR-11-LABX-0025-01) funded by the French program Investissement d’avenir. This work is published on *Computational Statistics*, doi:10.1007/s00180-017-0770-y.

## References

C. Andrieu and G. Roberts. The pseudo-marginal approach for efficient Monte Carlo computations. *The Annals of Statistics*, pages 697–725, 2009.

- C. Andrieu, A. Doucet, and R. Holenstein. Particle Markov chain Monte Carlo methods (with discussion). *Journal of the Royal Statistical Society: Series B*, 72(3):269–342, 2010.
- M. Beaumont. Estimation of population growth or decline in genetically monitored populations. *Genetics*, 164(3):1139–1160, 2003.
- O. Cappé, E. Moulines, and T. Rydén. *Inference in hidden Markov models*. Springer, 2005.
- O. Cappe, SJ Godsill, and E. Moulines. An overview of existing methods and recent advances in sequential Monte Carlo. *Proceedings of the IEEE*, 95:899–924, 2007.
- B. Delyon, M. Lavielle, and E. Moulines. Convergence of a stochastic approximation version of the EM algorithm. *Annals of Statistics*, pages 94–128, 1999.
- A.P. Dempster, N. Laird, and D.B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society, Series B*, 39(1):1–38, 1977.
- S. Ditlevsen and A. Samson. Estimation in the partially observed stochastic Morris-Lecar neuronal model with particle filter and stochastic approximation methods. *Annals of Applied Statistics*, 2:674–702, 2014.
- S. Donnet and A. Samson. Parametric inference for mixed models defined by stochastic differential equations. *ESAIM: Probability and Statistics*, 12:196–218, 2008.
- A. Doucet, N. de Freitas, and N. Gordon, editors. *Sequential Monte Carlo methods in practice*. Springer Science & Business Media, 2001.
- M. Fasiolo, N. Pya, and S. Wood. A comparison of inferential methods for highly nonlinear state space models in ecology and epidemiology. *Statistical Science*, 31(1):96–118, 2016.
- A. Golightly and D. Wilkinson. Bayesian parameter inference for stochastic biochemical network models using particle Markov chain Monte Carlo. *Interface focus*, 1(6):807–820, 2011.
- N. Gordon, D. Salmond, and A. Smith. Novel approach to nonlinear/non-Gaussian Bayesian state estimation. In *IEE Proceedings F-Radar and Signal Processing*, volume 140, pages 107–113, 1993.
- E Herbst and F Schorfheide. Tempered particle filtering. Technical Report 23448, National Bureau of Economic Research, May 2017.
- Q. Huys, M.B. Ahrens, and L. Paninski. Efficient estimation of detailed single-neuron models. *J Neurophysiol*, 96(2):872–890, 2006.
- Quentin J. M. Huys and Liam Paninski. Smoothing of, and Parameter Estimation from, Noisy Biophysical Recordings. *PLOS Computational Biology*, 5(5):e1000379, 2009.
- E. Ionides, D. Nguyen, Y. Atchadé, S. Stoev, and A. King. Inference for dynamic and latent variable models via iterated, perturbed Bayes maps. *Proceedings of the National Academy of Sciences*, 112(3):719–724, 2015.
- A. Jasra, S. Singh, J. Martin, and E. McCoy. Filtering via approximate Bayesian computation. *Statistics and Computing*, 22(6):1223–1237, 2012.
- N. Kantas, A. Doucet, S. Singh, J. Maciejowski, and N. Chopin. On particle methods for parameter estimation in state-space models. *Statistical Science*, 30(3):328–351, 2015.
- A. King, D. Nguyen, and E. Ionides. Statistical inference for partially observed Markov processes via the R package pomp. *Journal of Statistical Software*, 69(12), 2015.

- G. Kitagawa. Monte Carlo filter and smoother for non-Gaussian nonlinear state space models. *Journal of Computational and Graphical statistics*, 5(1):1–25, 1996.
- E. Kuhn and M. Lavielle. Maximum likelihood estimation in nonlinear mixed effects models. *Computational Statistics & Data Analysis*, 49(4):1020–1038, 2005.
- M. Lavielle. *Mixed effects models for the population approach: models, tasks, methods and tools*. CRC Press, 2014.
- F Lindsten. An efficient stochastic approximation EM algorithm using conditional particle filters. *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on*, pages 6274 – 6278, 2013.
- J. Liu. *Monte Carlo strategies in scientific computing*. Springer Science & Business Media, 2008.
- T. Louis. Finding the observed information matrix when using the EM algorithm. *Journal of the Royal Statistical Society. Series B*, pages 226–233, 1982.
- J. M. Marin, P. Pudlo, C. P. Robert, and R. Ryder. Approximate Bayesian computational methods. *Statistics and Computing*, 22(6):1167–1180, 2012.
- G Martin, B McCabe, D Frazier, W Maneesoonthorn, and CP Robert. Auxiliary likelihood-based approximate Bayesian computation in state space models. *arXiv:1604.07949*, 2016.
- O Papaspiliopoulos, G Roberts, and M Sköld. A general framework for the parametrization of hierarchical models. *Statistical Science*, pages 59–73, 2007.
- U. Picchini. Inference for SDE models via approximate Bayesian computation. *Journal of Computational and Graphical Statistics*, 23(4):1080–1100, 2014.
- U. Picchini and J. Forman. Accelerating inference for diffusions observed with measurement error and large sample sizes using approximate Bayesian computation. *Journal of Statistical Computation and Simulation*, 86:195–213, 2015.
- J. Pinheiro and D. Bates. Approximations to the log-likelihood function in the nonlinear mixed-effects model. *Journal of Computational and Graphical Statistics*, 4(1):12–35, 1995.
- C. Sherlock, A. Thiery, G. Roberts, and J. Rosenthal. On the efficiency of pseudo-marginal random walk Metropolis algorithms. *The Annals of Statistics*, 43(1):238–275, 2015.
- A Sitz, U Schwarz, J Kurths, and H Voss. Estimation of parameters and unobserved components for nonlinear systems from noisy time series. *Physical Review E*, 66(1):016210, 2002.
- D. Wilkinson. smfsb: Stochastic modelling for systems biology, 2015. <https://cran.r-project.org/web/packages/smfsb/>.

#### Conditional densities for the Gibbs sampler in section 4.1.4.

Here we report the conditional densities for the Gibbs sampler when  $\mathbf{X}$  is sampled in block. Here  $\pi(\sigma_x)$  and  $\pi(\sigma_y)$  are prior densities.

$$\begin{aligned}
p(\sigma_x | \sigma_y, \mathbf{X}, \mathbf{Y}) &\propto \prod_{j=1}^n \frac{1}{\sigma_x} e^{-\frac{1}{2\sigma_x^2} (X_j - 2 \sin(\exp(X_{j-1})))^2} \pi(\sigma_x) \\
p(\sigma_y | \sigma_x, \mathbf{X}, \mathbf{Y}) &\propto \prod_{j=1}^n \frac{1}{\sigma_y} e^{-\frac{1}{2\sigma_y^2} (Y_j - X_j)^2} \pi(\sigma_y) \\
p(\mathbf{X} | \sigma_x, \sigma_y, \mathbf{Y}) &\propto \prod_{j=1}^n \frac{1}{\sigma_x \sigma_y} e^{-\frac{1}{2\sigma_y^2} (Y_j - X_j)^2 - \frac{1}{2\sigma_x^2} (X_j - 2 \sin(\exp(X_{j-1})))^2}
\end{aligned}$$

## First and second derivatives for the example in section 4.1.

Here we report the first and second derivatives of the complete log-likelihood  $L_c(\mathbf{X}, \mathbf{Y}; \boldsymbol{\theta})$  with respect to  $\boldsymbol{\theta} = (\sigma_x^2, \sigma_y^2)$ .

$$\begin{aligned}
\frac{\partial L_c(\mathbf{X}, \mathbf{Y}; \boldsymbol{\theta})}{\partial \sigma_x^2} &= -\frac{n}{2\sigma_x^2} + \frac{1}{2\sigma_x^4} \sum_{j=1}^n [X_j - 2 \sin(\exp(X_{j-1}))]^2, \\
\frac{\partial L_c(\mathbf{X}, \mathbf{Y}; \boldsymbol{\theta})}{\partial \sigma_y^2} &= -\frac{n}{2\sigma_y^2} + \frac{1}{2\sigma_y^4} \sum_{j=1}^n (Y_j - X_j)^2, \\
\frac{\partial^2 L_c(\mathbf{X}, \mathbf{Y}; \boldsymbol{\theta})}{\partial \sigma_x^2 \partial \sigma_y^2} &= \frac{\partial^2 L_c(\mathbf{X}, \mathbf{Y}; \boldsymbol{\theta})}{\partial \sigma_y^2 \partial \sigma_x^2} = 0, \\
\frac{\partial^2 L_c(\mathbf{X}, \mathbf{Y}; \boldsymbol{\theta})}{\partial (\sigma_x^2)^2} &= \frac{n}{2\sigma_x^4} - \frac{1}{\sigma_x^6} \sum_{j=1}^n [X_j - 2 \sin(\exp(X_{j-1}))]^2, \\
\frac{\partial^2 L_c(\mathbf{X}, \mathbf{Y}; \boldsymbol{\theta})}{\partial (\sigma_y^2)^2} &= \frac{n}{2\sigma_y^4} - \frac{1}{\sigma_y^6} \sum_{j=1}^n (Y_j - X_j)^2.
\end{aligned}$$

## Fisher Information matrix for the example in section 4.2.

To compute the Fisher Information matrix as suggested in section 3 we need to differentiate the complete data log-likelihood with respect to the four parameters  $\boldsymbol{\theta} = (K_e, Cl, \sigma^2, \sigma_\varepsilon^2)$ . We differentiate w.r.t.  $(\sigma^2, \sigma_\varepsilon^2)$  instead of  $(\sigma, \sigma_\varepsilon)$  because the complete log-likelihood is expressed as a function of sufficient statistics for  $(\sigma^2, \sigma_\varepsilon^2)$ .

Set, for  $i = 1, \dots, N$ ,

$$z_i(\boldsymbol{\theta}) = x_i - x_{i-1} - h(Dose \cdot K_a \cdot \frac{K_e}{Cl} \cdot e^{-K_a \tau_{i-1}} - K_e \cdot x_{i-1}).$$

The four coordinates of the gradient are:

$$\begin{aligned}\frac{\partial}{\partial K_e} L_c(\mathbf{Y}, \mathbf{X}; \boldsymbol{\theta}) &= -\frac{1}{\sigma^2} \sum_{i=1}^N \frac{z_i(\boldsymbol{\theta})}{x_{i-1}} \left( x_{i-1} - \frac{Dose \cdot K_a}{Cl} e^{-K_a \tau_{i-1}} \right) \\ \frac{\partial}{\partial Cl} L_c(\mathbf{Y}, \mathbf{X}; \boldsymbol{\theta}) &= -\frac{1}{\sigma^2} \sum_{i=1}^N \frac{z_i(\boldsymbol{\theta})}{x_{i-1}} \left( \frac{Dose \cdot K_a \cdot K_e}{Cl^2} e^{-K_a \tau_{i-1}} \right) \\ \frac{\partial}{\partial \sigma^2} L_c(\mathbf{Y}, \mathbf{X}; \boldsymbol{\theta}) &= -\frac{N}{2\sigma^2} + \frac{1}{2h\sigma^4} \sum_{i=1}^N \frac{z_i(\boldsymbol{\theta})^2}{x_{i-1}} \\ \frac{\partial}{\partial \sigma_\varepsilon^2} L_c(\mathbf{Y}, \mathbf{X}; \boldsymbol{\theta}) &= -\frac{n}{2\sigma_\varepsilon^2} + \frac{1}{2\sigma_\varepsilon^4} \sum_{j=1}^n (y_j - x_j)^2.\end{aligned}$$

Entries for the Fisher information matrix are (recall this is a symmetric matrix, therefore redundant terms are not reported. Further missing entries consist of zeros):

$$\begin{aligned}\frac{\partial^2}{\partial^2 K_e} L_c(\mathbf{Y}, \mathbf{X}; \boldsymbol{\theta}) &= -\frac{h}{\sigma^2} \sum_{i=1}^N (x_{i-1} - Dose \cdot \frac{K_a}{Cl} \cdot e^{-K_a \tau_{i-1}})^2 \frac{1}{x_{i-1}} \\ \frac{\partial^2}{\partial^2 Cl} L_c(\mathbf{Y}, \mathbf{X}; \boldsymbol{\theta}) &= -\frac{1}{\sigma^2} \sum_{i=1}^N \left\{ \frac{1}{x_{i-1}} \left[ \frac{Dose \cdot K_a \cdot K_e}{Cl^2} e^{-K_a \tau_{i-1}} \left( h - \frac{2z_i(\boldsymbol{\theta})}{Cl} \right) \right] \right\} \\ \frac{\partial^2}{\partial^2 \sigma^2} L_c(\mathbf{Y}, \mathbf{X}; \boldsymbol{\theta}) &= \frac{N}{2\sigma^4} - \frac{1}{h\sigma^6} \sum_{i=1}^N \frac{z_i(\boldsymbol{\theta})^2}{x_{i-1}} \\ \frac{\partial^2}{\partial^2 \sigma_\varepsilon^2} L_c(\mathbf{Y}, \mathbf{X}; \boldsymbol{\theta}) &= \frac{n}{2\sigma_\varepsilon^4} - \frac{1}{\sigma_\varepsilon^6} \sum_{j=1}^n (y_j - x_j)^2\end{aligned}$$

$$\begin{aligned}\frac{\partial^2}{\partial K_e \partial Cl} L_c(\mathbf{Y}, \mathbf{X}; \boldsymbol{\theta}) &= -\frac{1}{\sigma^2} \sum_{i=1}^N \frac{1}{x_{i-1}} \left\{ \frac{Dose \cdot K_a}{Cl^2} e^{-K_a \tau_{i-1}} \left[ h \cdot K_e \left( x_{i-1} - \frac{Dose \cdot K_a}{Cl} e^{-K_a \tau_{i-1}} \right) \right. \right. \\ &\quad \left. \left. + z_i(\boldsymbol{\theta}) \right] \right\} \\ \frac{\partial^2}{\partial \sigma^2 \partial K_e} L_c(\mathbf{Y}, \mathbf{X}; \boldsymbol{\theta}) &= \frac{1}{\sigma^4} \sum_{i=1}^N \frac{z_i(\boldsymbol{\theta})}{x_{i-1}} \left( x_{i-1} - \frac{Dose \cdot K_a}{Cl} e^{-K_a \tau_{i-1}} \right) \\ \frac{\partial^2}{\partial \sigma^2 \partial Cl} L_c(\mathbf{Y}, \mathbf{X}; \boldsymbol{\theta}) &= \frac{1}{\sigma^4} \sum_{i=1}^N \frac{z_i(\boldsymbol{\theta})}{x_{i-1}} \left( \frac{Dose \cdot K_a \cdot K_e}{Cl^2} e^{-K_a \tau_{i-1}} \right).\end{aligned}$$