

SEQUENTIAL BAYESIAN INFERENCE FOR IMPLICIT HIDDEN MARKOV MODELS AND CURRENT LIMITATIONS

PIERRE E. JACOB¹

Dedicated to Philip Perry

Abstract. Hidden Markov models can describe time series arising in various fields of science, by treating the data as noisy measurements of an arbitrarily complex Markov process. Sequential Monte Carlo (SMC) methods have become standard tools to estimate the hidden Markov process given the observations and a fixed parameter value. We review some of the recent developments allowing the inclusion of parameter uncertainty as well as model uncertainty. The shortcomings of the currently available methodology are emphasised from an algorithmic complexity perspective. The statistical objects of interest for time series analysis are illustrated on a toy “Lotka-Volterra” model used in population ecology. Some open challenges are discussed regarding the scalability of the reviewed methodology to longer time series, higher-dimensional state spaces and more flexible models.

Résumé. Les modèles à chaîne de Markov cachée permettent de décrire les séries temporelles de divers domaines scientifiques, en traitant les données comme des mesures bruitées d’un processus de Markov arbitrairement complexe. Les méthodes de Monte Carlo séquentielles (SMC) sont devenues des outils standards pour l’estimation du processus de Markov caché sachant les observations et une valeur fixée du paramètre. Nous passons en revue quelques unes des récentes avancées permettant de prendre en compte l’incertitude sur le paramètre ainsi que l’incertitude sur le modèle. Les limites de la méthodologie actuelle sont discutées sous l’angle de la complexité algorithmique. Les objets statistiques d’intérêt pour l’analyse des séries temporelles sont illustrés sur un modèle jouet de type “Lotka-Volterra” utilisé en écologie des populations. Quelques questions ouvertes sont finalement posées concernant l’extension de la méthodologie présentée pour traiter des séries de données plus longues, des espaces d’états de dimension plus grande et des modèles plus flexibles.

1. SETTING

1.1. Hidden Markov models

Hidden Markov models constitute a very flexible class of models for time series data. Consider time series made of real-valued vectors $y_t \in \mathbb{Y} \subset \mathbb{R}^{d_y}$ for a countable collection of times $t \in \{t_1, \dots, t_n, \dots\}$ and some integer d_y . For simplicity we consider integer-valued times. Hidden Markov models propose to treat the observations $(y_t)_{t \in \mathbb{N}}$ as if they were arising from noisy measurements of a latent Markov chain $(x_t)_{t \in \mathbb{N}}$. First, the distribution of a Markov chain $(x_t)_{t \in \mathbb{N}}$ living in $\mathbb{X} \subset \mathbb{R}^{d_x}$, for some integer d_x , is specified through the distribution $\mu(dx_0)$ of its initial state and the conditional distribution of each successive state $f(dx_t | x_{t-1})$, which is called the transition

¹ Department of Statistics, University of Oxford, pierre.jacob@stats.ox.ac.uk.




FIGURE 1. Graphical representation of the variables defined by a hidden Markov model \mathcal{M} .




FIGURE 2. A time series of 365 observations generated according to the phytoplankton–zooplankton model described in Section 1.3. The observations represent daily measurements of phytoplankton concentrations in a volume of water.

distribution. Then the model specifies the distribution $g(dy_t | x_t)$ of each observation given the current hidden state, which is called the measurement distribution. Finally all those distributions are parametrized by a vector θ living in a set $\Theta \subset \mathbb{R}^{d_\theta}$ for some integer d_θ . We explicitly write the parameter θ in $\mu(dx_0 | \theta)$, $f(dx_t | x_{t-1}, \theta)$ and $g(dy_t | x_t, \theta)$, and a model \mathcal{M} refers to the collection of objects $\{\Theta, \mu, f, g\}$. In the following, for a sequence $(v_t)_{t \in \mathbb{N}}$ (resp. $(v^t)_{t \in \mathbb{N}}$) and $0 \leq s \leq t$, the vector (v_s, \dots, v_t) (resp. (v^s, \dots, v^t)) is denoted by $v_{s:t}$ (resp. $v^{s:t}$).

Hidden Markov models are often represented by dependency graphs such as Figure 1, although θ is typically omitted. The graph indicates that the law of y_t given x_t and θ is independent of the other variables, the law of x_t given x_{t-1}, y_t, x_{t+1} and θ is independent of the other variables, and so on. Figure 2 shows a time series generated by a hidden Markov model to be specified in Section 1.3.

After having obtained some data, the user comes up with one or several candidate models, denoted by $\mathcal{M}^{(1)}, \dots, \mathcal{M}^{(M)}$, for some integer M . Hidden Markov models are routinely applied to the modeling of volatility of financial assets, of meteorological time series such as wind speed and direction, or of population growth in ecology (see [18, 40] for a variety of applications). The design of a model preferably takes into account as much knowledge as possible on the phenomenon under study, and a particular model developed for one phenomenon is unlikely to be of use for others. There are exceptions: for instance, stochastic volatility models have been applied to pollution data in [1].

In the rest of the article we assume that a collection of models is given to us. Then the following questions naturally arise.

- For each model, how do the data inform about the parameters?
- For each model, how do the data inform about the latent Markov process?
- How do the data inform about the choice of a model?
- How to predict future observations?

If one can simulate synthetic datasets from a model given any parameter value, then by trial and error, one can gather some intuition on how models perform. Intuitively, a model that generates synthetic time series that resemble actual data, at least for some parameter value, is believed to be a good model, for instance in the sense that one would hope its prediction of future observations to be reliable.

Statistical inference is concerned with formalizing and automatizing this ad hoc procedure, by making a principled connection between data and models. In Section 1.2 the Bayesian framework is shown to transform the above questions into various integrals. In Section 1.3 we introduce the notion of implicit model, on which we focus thereafter. In Section 1.4 we state some desired properties of statistical computing methods for time series models. Section 2 reviews some numerical methods compatible with implicit models, called “plug and play” methods, and discusses whether they meet the desired requirements. In particular, sequential Monte Carlo (SMC) and particle Markov chain Monte Carlo methods are introduced. In Section 3 we describe a recently proposed method called SMC² to perform Bayesian inference sequentially and mention its shortcomings. Section 4 illustrates the method on a population growth model. Finally Section 5 discusses the reviewed methodology and some open challenges.

1.2. Objects of inference

To begin with, the model \mathcal{M} and a particular parameter value θ are considered fixed, and the goal is to estimate the distribution of the hidden process given observations. Filtering refers to the task of estimating at time t the current state x_t given the available observations $y_{0:t}$. The filtering distribution $p(dx_t | y_{0:t}, \theta)$ is obtained using Bayes rule as

$$\int_{\mathbb{X}} \varphi(x_t) p(dx_t | y_{0:t}, \theta) = \frac{1}{p(y_{0:t} | \theta)} \int_{\mathbb{X}^{t+1}} \varphi(x_t) \underbrace{\left(\mu(dx_0 | \theta) \prod_{s=1}^t f(dx_s | x_{s-1}, \theta) \right)}_{\text{law of the latent Markov chain}} \underbrace{\left(\prod_{s=0}^t g(y_s | x_s, \theta) \right)}_{\text{conditional law of the observations}} \quad (1)$$

for any test function φ . In the rest of the article, φ refers to a generic test function, defined either on \mathbb{X} , \mathbb{Y} or Θ . The normalizing constant $p(y_{0:t} | \theta)$ in Eq. (1) is called the marginal likelihood of θ and we come back to it in Eq. (4). Filtering sometimes refers to the distribution $p(dx_{0:t} | y_{0:t}, \theta)$ of the path (or “trajectory”) $x_{0:t}$ given $y_{0:t}$. Prediction refers to the inference of both a future observation y_{t+k} and a future state x_{t+k} given the current observations $y_{0:t}$, for some $k \geq 1$. Whether the interest lies on future states or future observations depends on the application. Using the Markov property, prediction is obtained as a by-product of the filtering distribution. Denoting the state predictive distribution by $p(dx_{t+k} | y_{0:t}, \theta)$, it can be written as

$$\int_{\mathbb{X}} \varphi(x_{t+k}) p(dx_{t+k} | y_{0:t}, \theta) = \int_{\mathbb{X}^{k+1}} \varphi(x_{t+k}) \underbrace{p(dx_t | y_{0:t}, \theta)}_{\text{filtering law}} \underbrace{\left(\prod_{s=t+1}^{t+k} f(dx_s | x_{s-1}, \theta) \right)}_{\text{conditional law of the future Markov chain}} \quad (2)$$

The predictive distribution of y_{t+k} given $y_{0:t}$, denoted by $p(dy_{t+k} | y_{0:t}, \theta)$, is then derived as

$$\int_{\mathbb{Y}} \varphi(y_{t+k}) p(dy_{t+k} | y_{0:t}, \theta) = \int_{\mathbb{Y}} \int_{\mathbb{X}} \varphi(y_{t+k}) \underbrace{p(dx_{t+k} | y_{0:t}, \theta)}_{\text{state prediction}} \underbrace{g(dy_{t+k} | x_{t+k}, \theta)}_{\text{measurement}} \quad (3)$$

Finally smoothing refers to inference of a past state x_{t-k} given $y_{0:t}$, for some $1 \leq k \leq t$.

In many realistic situations one does not know which parameter value θ to plug in the model, even if the interest lies in filtering the hidden process; in other settings parameters are themselves the objects of interest. To learn about the parameters from the observations, a “prior” probability distribution π_θ is given to the parameter θ and the goal is to estimate the “posterior” distribution given the observations. The marginal likelihood of θ is the probability density function of the distribution of the data given θ , evaluated at the observations $y_{0:t}$:

$$p(y_{0:t} | \theta) = \int_{\mathbb{X}^{t+1}} \mu(dx_0 | \theta) \prod_{s=1}^t f(dx_s | x_{s-1}, \theta) \prod_{s=0}^t g(y_s | x_s, \theta). \quad (4)$$

The marginal likelihood of θ was the normalizing constant in Eq. (1); the “full likelihood” sometimes refers to the joint probability density function of $x_{0:t}$ and $y_{0:t}$ given θ . Then the posterior distribution $\pi_{\theta,t}$ of θ given $y_{0:t}$ is defined by Bayes rule as

$$\int_{\Theta} \varphi(\theta) \pi_{\theta,t}(d\theta) = \frac{1}{\int_{\Theta} p(y_{0:t} | \theta) \pi_\theta(d\theta)} \int_{\Theta} \varphi(\theta) p(y_{0:t} | \theta) \pi_\theta(d\theta). \quad (5)$$

The normalizing constant in Eq. (5) will be useful for model comparison, and we come back to it in Eq. (7). By taking parameter uncertainty into account, one can redefine the tasks of filtering, prediction and smoothing. For instance, filtering under parameter uncertainty refers to the distribution $p(dx_t | y_{0:t})$ of the current hidden state x_t given $y_{0:t}$, averaged over all possible parameters:

$$\int_{\mathbb{X}} \varphi(x_t) p(dx_t | y_{0:t}) = \int_{\Theta} \int_{\mathbb{X}} \varphi(x_t) \underbrace{p(dx_t | y_{0:t}, \theta)}_{\text{filtering given parameter}} \underbrace{\pi_{\theta,t}(d\theta)}_{\text{posterior on parameter}}. \quad (6)$$

Likewise one can be interested in filtering over the full paths, prediction and smoothing under parameter uncertainty.

Once parameter uncertainty is taken into account, the next source of uncertainty is at the model level. When several models $\mathcal{M}^{(1)}, \dots, \mathcal{M}^{(M)}$ are available, there are various ways to use the observations to compare models (see Chapter 7 of [75], Chapter 6 of [11]). A building block of model comparison is the “model evidence”. The evidence of $\mathcal{M}^{(m)}$ is defined as the normalizing constant $\mathcal{Z}^{(m)}$ of its posterior distribution, that is, denoting the parameter $\theta^{(m)}$ and its space as $\Theta^{(m)}$,

$$\mathcal{Z}_t^{(m)} = \int_{\Theta^{(m)}} p(y_{0:t} | \theta) \pi_{\theta^{(m)}}(d\theta) = p(y_{0:t} | \mathcal{M}^{(m)}). \quad (7)$$

The model evidence, which was the normalizing constant in Eq. (5), can be understood as the density of the observations $y_{0:t}$ given the model $\mathcal{M}^{(m)}$. By introducing prior probabilities on the discrete set of model labels $\{\mathcal{M}^{(1)}, \dots, \mathcal{M}^{(M)}\}$, one can then consider posterior probabilities of the models given the data, obtained again by Bayes rule as

$$\mathbb{P}(\mathcal{M} = \mathcal{M}^{(m)} | y_{0:t}) = \frac{\mathbb{P}(\mathcal{M} = \mathcal{M}^{(m)}) \mathcal{Z}_t^{(m)}}{\sum_{m'=1}^M \mathbb{P}(\mathcal{M} = \mathcal{M}^{(m')}) \mathcal{Z}_t^{(m')}}. \quad (8)$$

On top of parameter uncertainty, model uncertainty can be taken into account when performing filtering, prediction or smoothing. For instance the predictive distribution $\mathcal{P}(dy_{t+k} | y_{0:t})$ of new observations y_{t+k} given

$y_{0:t}$ under both model and parameter uncertainty can be written as

$$\begin{aligned} & \int_{\mathbb{Y}} \varphi(y_{t+k}) \mathcal{P}(dy_{t+k} \mid y_{0:t}) \\ &= \sum_{m=1}^M \int_{\Theta^{(m)}} \int_{\mathbb{Y}} \varphi(y_{t+k}) \underbrace{p(dy_{t+k} \mid y_{0:t}, \theta, \mathcal{M}^{(m)})}_{\text{predictive distribution of } y} \underbrace{\pi_{\theta^{(m)},t}(d\theta)}_{\text{posterior on parameter}} \underbrace{\mathbb{P}(\mathcal{M} = \mathcal{M}^{(m)} \mid y_{0:t})}_{\text{posterior on model}}. \end{aligned} \quad (9)$$

This type of quantity is also referred to as model averaging [52]. Sometimes the task consists in selecting one model among the M proposed ones. Then a standard procedure is to estimate the posterior odds of model $\mathcal{M}^{(m)}$ versus model $\mathcal{M}^{(m')}$:

$$\frac{\mathbb{P}(\mathcal{M} = \mathcal{M}^{(m)} \mid y_{0:t})}{\mathbb{P}(\mathcal{M} = \mathcal{M}^{(m')} \mid y_{0:t})} = \underbrace{\frac{p(y_{0:t} \mid \mathcal{M}^{(m)})}{p(y_{0:t} \mid \mathcal{M}^{(m')})}}_{\text{Bayes factor}} \times \underbrace{\frac{\mathbb{P}(\mathcal{M} = \mathcal{M}^{(m)})}{\mathbb{P}(\mathcal{M} = \mathcal{M}^{(m')})}}_{\text{prior odds}}. \quad (10)$$

Typically the value 1 is assigned to the prior odds, corresponding to uniform prior probabilities on the model labels, and thus the posterior odds correspond to the Bayes factor [61]. It is well-known [66] that the Bayes factor embodies the principle of Occam's razor: "simple" models (i.e. with a low-dimensional parameter space) are favoured over "complex" models (i.e. with a high-dimensional parameter space) until enough data have accrued to support the additional complexity.

As discussed in [11], it can be confusing to specify non-zero prior probabilities on the event "model $\mathcal{M}^{(m)}$ is the true, data-generating model", for each $m \in \{1, \dots, M\}$. Often, we do not believe that any of the M candidate models is the true data-generating model. In that case, we can bypass the specification of probabilities of the models being true, and interpret the model evidence in Eq. (7) as a prior predictive criterion [31], representing how likely the observations are under the model. The logarithm of the Bayes factor is then the difference between expected utilities associated to each model, under the logarithmic scoring rule, for the task of predicting observations using the prior [61]. Thus, even in the setting where each candidate model is mis-specified compared to the data-generating process, the use of Bayes factors is still defensible.

To summarise, there are three layers of uncertainty in hidden Markov models: the hidden process, the model parameters and the model itself. Ideally each of these uncertainties should be taken into account.

1.3. Implicit models

The integrals presented above are in general impossible to evaluate, except for linear Gaussian models. In these models the latent Markov chain is an autoregressive process and the observations are Gaussian measurements of it. Formally $x_0 \sim \mathcal{N}(\mu_0, \Sigma_0)$, the Gaussian distribution with mean μ_0 and variance Σ_0 , and for each $t \geq 1$,

$$x_t = Ax_{t-1} + v_t \quad \text{and} \quad y_{t-1} = Bx_{t-1} + w_{t-1},$$

where v_t and w_t are Gaussian random variables with zero mean and variances Σ_x and Σ_y respectively. The parameter is thus $\theta = (\mu_0, \Sigma_0, A, \Sigma_x, B, \Sigma_y)$ and is made either of real values or of vectors and matrices of compatible dimensions. The linearity and Gaussianity of the model equations imply that various conditional distributions of interest, for instance the filtering distribution of x_t given $y_{0:t}$ and θ , are also Gaussian. The Kalman filter [2] provides the mean and variance of these Gaussian distributions for a computational cost of $\mathcal{O}(t)$. As a by-product the likelihood in Eq. (4) can be evaluated for any θ for a cost in $\mathcal{O}(t)$, which allows parameter estimation using standard techniques, sometimes under the name of system parameter identification, see e.g. [82]. In a nutshell, linearity and Gaussianity make integration with respect to the hidden Markov chain

analytically possible and thus the integrals of Section 1.2 can be either evaluated or at least approximated using standard techniques such as Markov chain Monte Carlo.

As discussed in [71] there are some non-Gaussian models for which the filtering distribution is tractable. However it has been argued, see e.g. [17], that models should preferably be proposed based on scientific grounds rather than on computational ones. In practice scientists often come up with complex and generative models, and then use linearization and Gaussian assumptions only to enjoy the numerical efficiency of Kalman filters and related methods (e.g. Extended Kalman filter, Ensemble Kalman filter), particularly in high-dimensional settings [45]. Linearization and Gaussianity assumptions result in a systematic bias in the subsequent estimation, compared to the results that would be obtained under the original model.

Let us have a look at a simple generative model that has been proposed to model the interaction of phytoplankton and zooplankton in [58, 69]. The ‘‘PZ’’ model is a variation of a Lotka-Volterra model for interactions between prey and predator. Phytoplankton are modeled as prey on which zooplankton are grazing. Over successive days indexed by t , the model describes the stochastic growth rate of prey α_t , the population size of prey p_t and the population size of predator z_t . Thus the hidden state $x_t = (\alpha_t, p_t, z_t)$ is three-dimensional. The stochastic growth rate α_t follows the same distribution every day $t \geq 0$: $\alpha_t \sim \mathcal{N}(\mu_\alpha, \sigma_\alpha^2)$. The initial distributions for both species are given by

$$\log p_0 \sim \mathcal{N}(\mu_p, \sigma_p^2) \quad \text{and} \quad \log z_0 \sim \mathcal{N}(\mu_z, \sigma_z^2).$$

The transition of (p_t) and (z_t) is jointly described by the differential equation

$$\begin{aligned} \frac{dp_t}{dt} &= \alpha p_t - c p_t z_t, \\ \frac{dz_t}{dt} &= e c p_t z_t - m_l z_t - m_q z_t^2. \end{aligned}$$

This is to be interpreted as follows: given a value for p_{t-1} and z_{t-1} , and a draw α of α_t , the next states p_t and z_t are obtained as the deterministic solution of the above equation over one time unit. In the equation, c represents the clearance rate of the prey, e is the growth efficiency of the predator, m_l and m_q are the linear and quadratic mortality rates of the predator. Note that (p_t, z_t) can also be defined at non-integer times, and we could consider α_t to be a piecewise constant continuous time process jumping at each integer time according to $\mathcal{N}(\mu_\alpha, \sigma_\alpha^2)$. However the observations are gathered in discrete time and thus we find it more convenient to specify the hidden Markov chain in discrete time as well. To summarize, given $x_{t-1} = (\alpha_{t-1}, p_{t-1}, z_{t-1})$, the next state is obtained by drawing α_t from $\mathcal{N}(\mu_\alpha, \sigma_\alpha^2)$ and (p_t, z_t) is the solution of the ordinary differential equation above. In practice, this solution can be approximated with arbitrary precision by numerical solvers such as Runge-Kutta methods. Note also that the difference between the PZ model and the classical Lotka-Volterra lies in the addition of a quadratic mortality term, and in the randomness of the growth rate α_t sampled at each integer time.

Given the Markov process $(x_t)_{t \in \mathbb{N}}$, the observations are noisy measurement of the phytoplankton, $\log y_t \sim \mathcal{N}(\log p_t, \sigma_y^2)$; the zooplankton are not measured. Indeed it is comparatively easier to measure the concentration of phytoplankton in a volume of water due to the fluorescence of the chlorophyll that they contain. For simplicity we set $c = 0.25$ and $e = 0.3$, $\mu_p = \mu_z = \log 2$, $\sigma_p = 0.2$, $\sigma_z = 0.1$, and leave the remaining constants as unknown (or ‘‘free’’) parameters. Thus we introduce $\theta = (\mu_\alpha, \sigma_\alpha, \sigma_y, m_l, m_q)$. Figure 2 was obtained by generating a year of data from the model, with parameters $\mu_\alpha = 0.7, \sigma_\alpha = 0.5, \sigma_y = 0.2, m_l = 0.1, m_q = 0.1$. A prior distribution is put on the model parameters, simply chosen to be a uniform distribution on $[0, 1]$ for each of the five components. For this application the interest lies both in the prediction of future states and in parameter inference.

The PZ model is very standard from a scientific point of view as Lotka-Volterra equations date back to the 1930s (see the historical introduction in [12]). However, from a statistical point of view the model is not linear nor Gaussian and thus the integrals defined in Section 1.2 are impossible to evaluate exactly. The model is generative in the sense that trajectories $x_{0:t}$ of the hidden Markov chain can be sampled, if not exactly, at least

with arbitrary precision using numerical solvers of ordinary differential equations, for any parameter θ . Then a series of observations $y_{0:t}$ can be simulated given a path $x_{0:t}$ and θ . When the transition distribution f is such that x_t can be sampled given x_{t-1} and θ but its transition density cannot be evaluated, the model is said to be “implicit” [17]. Here the ability to evaluate the transition density $f(x_t | x_{t-1}, \theta)$ for a given triplet (x_t, x_{t-1}, θ) would in fact depend on the numerical solver being used, but in general we do not want to assume that we know how to perform this computation.

Aside from Markov chains defined by differential equations, another scenario where sampling from the transition distribution is easier than evaluating its probability density function occurs when the transition involves latent variables. An example is the Lévy driven stochastic volatility model described in [8]. Given the previous state x_{t-1} and a parameter θ , x_t is obtained by sampling an integer-valued random variable k , and then k other random variables $v_{1:k}$ independently from some distribution $p(dv | \theta)$. The state x_t is obtained as $x_t = \psi(x_{t-1}, k, v_{1:k})$ for some deterministic function ψ that can be evaluated point-wise. Thus x_t is straightforward to simulate, but the evaluation of its transition density involves an integral of ψ over k and $v_{1:k}$, which is not analytically available for general functions ψ and random variables $(k, v_{1:k})$.

1.4. Online and exact inference

As mentioned earlier, the integrals of Section 1.2 are impossible to evaluate exactly for general hidden Markov models such as the PZ model described in Section 1.3. Let us generically denote by I_t one of these integrals, for instance the one in Eq. (6). Monte Carlo methods have been actively developed for hidden Markov models inside and outside the Bayesian framework [18, 76], and yield a random variable \hat{I}_t estimating I_t . Let us list some desirable features of \hat{I}_t and of the algorithm producing it, in the context of time series.

Numerical methods are said to be “exact” if, at any time t they produce consistent estimators of the integral I_t , where consistency is with respect to a tuning parameter N of the algorithm producing the estimator. For instance, if \hat{I}_t converges to I_t when N goes to infinity in the L_2 or “mean square” sense, then both the bias and the variance of the estimator go to zero, and the algorithm is considered exact. On the contrary the use of Extended Kalman filters for non-linear non-Gaussian models results in a systematic bias that cannot be reduced to zero, and this bias is typically hard to quantify. More trivially, the estimator that always returns “one” has zero variance, a bias that is uniformly bounded over the time index t (if I_t is so itself); but it does not have an algorithmic parameter N allowing a trade-off between computational power and precision.

Numerical methods for time series are “sequential” if an already obtained estimator can be “updated” upon the arrival of a new observation. For instance if an estimator \hat{I}_t of I_t has been obtained at time t , then a sequential method yields \hat{I}_{t+1} once y_{t+1} is made available, for a computational cost independent of t . The advantage of sequential methods for time series will be illustrated in Section 4. A sequential method is said to be “online” only if its performance does not deteriorate with t . Introducing the relative error $r(\hat{I}_t)$ of the estimator, typically defined as

$$r(\hat{I}_t) = \frac{\left(\mathbb{E} \left[(\hat{I}_t - I_t)^2 \right]\right)^{1/2}}{|I_t|},$$

then the method is online if $r(\hat{I}_t)$ is uniformly bounded from above over the time index t . This requirement rules out the use of standard Monte Carlo algorithms for hidden Markov models. For instance, in the case of Eq. (6), a trivial sequential importance sampling estimator would sample N draws from the “prior” distribution $\pi_\theta(d\theta)p(dx_{0:t} | \theta)$ and use the conditional density $p(y_{0:t} | x_{0:t}, \theta)$ as an importance weight. The resulting estimator can be updated for a fixed cost per observation and yields a consistent answer when N goes to infinity. However its variance would typically grow exponentially with t , and thus sequential importance sampling is not online for this problem.

In the next section we review existing sequential and exact Monte Carlo methods to approximate the objects described in Section 1.2, and we discuss whether or not they satisfy the above online requirement.

2. PLUG AND PLAY METHODS

Because of the numerous examples of implicit models as the PZ model of Section 1.3 we focus on numerical methods that are compatible with implicit models. These methods, called “plug and play” (or “equation free”) in [17], require only the ability to sample the hidden Markov chain, and to evaluate the measurement density.

2.1. Approximate Bayesian Computation

Perhaps the most natural plug and play algorithm for implicit models is the ABC (Approximate Bayesian Computation) method [9, 32, 85]. In a nutshell, ABC draws approximately from the posterior distribution of $(\theta, x_{0:t})$ given $y_{0:t}$ using the following steps.

- (1) Draw θ from the prior distribution π_θ .
- (2) Draw $x_{0:t}$, a realisation of the hidden Markov chain given θ .
- (3) Draw $\hat{y}_{0:t}$, a realisation of the observations given $x_{0:t}$ and θ .
- (4) If $\mathcal{D}(\hat{y}_{0:t}, y_{0:t}) \leq \varepsilon$, retain $(\theta, x_{0:t})$, otherwise go back to step (1).

In this algorithm, \mathcal{D} can be understood loosely as a distance defined on the observation space \mathbb{Y} [67]. The value of ε has to be chosen by the user; for smaller values, the synthetic data has to be closer (in the sense of \mathcal{D}) to the true data in order for $(\theta, x_{0:t})$ to be retained. It can be shown that if \mathcal{D} is a true distance on \mathbb{Y} then when ε goes to zero, the procedure samples from the true posterior distribution. When \mathcal{D} is not a distance (for instance when it is based on summary statistics), or when ε is a fixed value, then the samples obtained from ABC are not distributed according to the posterior distribution, and it is notoriously difficult to quantify the bias between the obtained approximation and the target distribution. Thus ABC estimators are typically not exact in the sense of Section 1.4. Note that ABC only requires the ability to sample from μ , f and g .

The next section introduces the particle filter, which is an exact, online and plug and play method to perform filtering and prediction for a given parameter value. On the other hand, particle filters require point-wise evaluations of the measurement density g , therefore they are less generally applicable than ABC.

2.2. Particle filters

Particle filters have become the preferred methods to deal with filtering, prediction and smoothing tasks [35, 39], since the publication of the seminal paper [51]. They were initially introduced to evaluate integrals such as Eq. (1) for non-linear, non-Gaussian hidden Markov models. The original algorithm, or “bootstrap particle filter”, is described in Algorithm 1.

Algorithm 1 Particle filter for a given parameter θ .

- 1: Draw for each $k \in \{1, \dots, N_x\}$ $x_0^k \sim \mu(dx_0 | \theta)$.
 - 2: **for** $t = 0$ to T **do**
 - 3: [weighting] Compute for each k , $w_t^k = g(y_t | x_t^k, \theta)$.
 - 4: [resampling] Sample ancestors $a_t^{1:N_x} \sim r(da^{1:N_x} | w_t^{1:N_x})$.
 - 5: [transition] Draw for each k , $x_{t+1}^k \sim f(dx_{t+1} | x_t^k, \theta)$.
 - 6: **end for**
-

The algorithm requires the user to specify a number of “particles” $N_x \in \mathbb{N}$. The transition and weighting steps propagate the samples from one distribution to the next following standard importance sampling [76]. The resampling step consists in selecting the particles according to their weights, so that the particles with lowest weights are killed while the particles with highest weights get replicated and propagated. Without the resampling step, only one particle would have a significant weight after a few time steps. Various resampling schemes exist as described in [39]. A standard resampling scheme consists in drawing each ancestor a_t^k independently from a categorical distribution with parameters $w_t^{1:N_x}$, so that for each $i \in \{1, \dots, N_x\}$, $\mathbb{P}(a_t^k = i) = w_t^i / \sum_l w_t^l$. In Algorithm 1, the resampling distribution is denoted by $r(da^{1:N_x} | w_t^{1:N_x})$ on Line 4.

The algorithm yields at each step t a weighted sample $(x_t^k, w_t^k)_{k=1}^{N_x}$ approximating the filtering distribution $p(dx_t | y_{0:t}, \theta)$, in the sense that the integral in Eq. (1) can be approximated by the weighted average

$$\frac{1}{\sum_{k=1}^{N_x} w_t^k} \sum_{k=1}^{N_x} w_t^k \varphi(x_t^k) \quad (11)$$

in a consistent manner as N_x goes to infinity. A rich theoretical literature supports particle filters. Some of the important results state that the estimator in Eq. (11) satisfies a central limit theorem (CLT), that both the bias and the variance are of order $1/N_x$, and that for a finite number N_x , the variance of the estimator can be uniformly bounded over the time steps t [26, 33, 35, 89]. The CLT and the time uniform results are remarkable since they make the particle filter “online” and “exact” in the sense of Section 1.4. Particle filters have therefore become standard tools for filtering in hidden Markov models. Kalman filter techniques such as Ensemble Kalman Filters are still used when the dimension d_x of the state space \mathbb{X} is large [45], because the variance of the particle filter estimators typically grows exponentially with d_x . Algorithmic improvements and theoretical studies of the impact of the dimension on particle filters have been recently proposed in [13, 15, 74]. Exact and online methods for large dimensional filtering problems constitute an active area of research.

An important by-product of this algorithm is an estimator $\widehat{Z}_t(\theta)$ of the marginal likelihood of θ at time t , as defined in Eq. (4). The estimator takes the simple form

$$\widehat{Z}_t(\theta) = \prod_{s=0}^t \left(\frac{1}{N_x} \sum_{k=1}^{N_x} w_s^k \right), \quad (12)$$

and can thus be updated sequentially at each step of Algorithm 1. This estimator has also been extensively studied in the literature. It happens to be unbiased, and by scaling the number of particles N_x linearly with the number of observations t , its relative variance is bounded by a constant, as proved in [22]. Hence the estimation of the likelihood using particle filters is not “online”: for a fixed cost per observation, the relative error goes to infinity. The cost of estimating the likelihood $p(y_{0:t} | \theta)$ can be said to be quadratic in the sense that one needs to choose $N_x = t$ to guarantee a fixed relative error. Then the cost of running a particle filter with N_x particles for t steps is t^2 , in the number of evaluations of g and draws from f .

To perform filtering on the path space or smoothing, one can simply modify Algorithm 1 to keep track of the generated paths $x_{0:t}^k$ instead of the most recent states x_t^k , as in Section 3.5 of [39]. Thus one would define $\bar{x}_0^k = x_0^k$ on Line 1 of Algorithm 1, and then on Line 5 one would define $\bar{x}_{0:t+1}^k = (\bar{x}_{0:t}^k, x_{t+1}^k)$. The resulting “path particles” $(\bar{x}_{0:t}^k, w_t^k)_{k=1}^{N_x}$ approximate the filtering distribution on the full path $p(dx_{0:t} | y_{0:t}, \theta)$, mentioned in Section 1.2. Thus, the estimator

$$\frac{1}{\sum_{k=1}^{N_x} w_t^k} \sum_{k=1}^{N_x} w_t^k \varphi(\bar{x}_{0:t}^k)$$

converges to the integral of φ with respect to $p(dx_{0:t} | y_{0:t}, \theta)$ when N_x goes to infinity and satisfies a CLT (Chapter 9 of [33]). However when t increases the variance of this estimator quickly deteriorates due to the well-known path degeneracy phenomenon. The variance has been shown to increase at least quadratically, and in general exponentially, as a function of t in [34, 41]. Indeed the resampling steps prune the population of distinct path particles at each time step. Let us denote by $\bar{x}_{0:t}^k(s)$, for $0 \leq s \leq t$, the s -th component of a path $\bar{x}_{0:t}^k$. Then at a given time t , the latest components $\bar{x}_{0:t}^{1:N_x}(t)$ of the path particles $\bar{x}_{0:t}^{1:N_x}$ are all distinct, but the first components $\bar{x}_{0:t}^{1:N_x}(0)$ contain many replicate values. In fact the number of unique values in $\bar{x}_{0:t}^{1:N_x}(0)$ quickly decreases to only one as t increases. More precisely the number of unique elements among the $N_x \times (t+1)$ components that compose the N_x path particles $\bar{x}_{0:t}^{1:N_x}$ has been upper bounded by $(t+1) + C N_x \log N_x$ in expectation in [56], where C is independent of t and N_x . To resolve the path degeneracy issue for the problem of smoothing given a parameter value θ , many particle algorithms have been proposed, such as fixed-lag

approximations, forward filtering backward smoothing or two filter formula, as described in [38, 39]. However the path degeneracy phenomenon also has consequences on parameter estimation, as described in the next section.

2.3. Particle-based approaches to parameter estimation

The early attempts to estimate the parameters using particle methods involve a reparametrization where the parameters θ are treated as an extra component of the hidden states. Thus a new hidden Markov model is introduced, where the hidden state is $\tilde{x}_t = (x_t, \theta_t)$ for all times t , x_t being the hidden state of the original model. The new initial distribution is then $\pi_\theta(d\theta_0)\mu(dx_0 | \theta_0)$, where π_θ is the prior on the parameters of the original model, and μ the original initial distribution. The new transition is $\delta_{\theta_{t-1}}(d\theta_t)f(dx_t | x_{t-1}, \theta_t)$, where δ_x represents the Dirac measure centered at the point x and f is the transition of original model. Finally the new measurement distribution is defined as $g(dy_t | x_t, \theta_t)$. Then by performing filtering on the modified model, one obtains a particle approximation of the distribution of \tilde{x}_t given $y_{0:t}$, and the distribution $\pi_{\theta,t}(d\theta)$ of θ given $y_{0:t}$ is obtained as a marginal distribution thereof.

The idea traces back to [62], who already recognized the occurrence of path degeneracy. Indeed, the parameter values $\theta_t^{1:N_x}$ are resampled as part of the states, but contrary to the states $x_t^{1:N_x}$ they are never diversified for the transition of the parameters is a delta function. Hence there are fewer and fewer unique values in the particle approximation of the posterior distribution $\pi_{\theta,t}(d\theta)$. We recognize the similarity with particle filtering on the path space, as described in the previous section. Early attempts such as [62, 65] proposed to replace the delta function by a Gaussian random walk in order to introduce diversity among the parameter samples. Alternatively, it has been proposed to introduce Markov chain Monte Carlo (MCMC) moves to diversify the parameter values [43, 50]. Those moves have the benefit of leaving the correct posterior distribution invariant. However the high dimensionality of $(\theta, x_{0:t})$ makes the design of efficient Markov chain Monte Carlo moves challenging in the setting of hidden Markov models, as well as the high correlation between the parameters and the states and between consecutive states [72]. Finally [43, 83] proposed specific moves in models such that the distribution of θ given $x_{0:t}$ and $y_{0:t}$ only depends on $x_{0:t}$ through a low dimensional sufficient statistic; see the early criticism in [5]. Reviews of various parameter estimation methods are proposed in [59, 60].

The inefficiency of standard MCMC moves and the path degeneracy phenomenon have made the various attempts at estimating the parameters as part of the hidden states generally unsuccessful. In the recent years two major advances have been proposed based on particle filters. The first one is Iterated filtering [54, 55], which is an optimization method relying on particle filters and an original representation of the score to find the maximum likelihood estimator in implicit models. The second one is particle Markov chain Monte Carlo [6], a class of MCMC algorithms using particle filters to design efficient proposal distributions on the space of $(\theta, x_{0:t})$. Here we recall a particle MCMC method called particle marginal Metropolis-Hastings (PMMH).

The pseudo-code is given in Algorithm 2. As described in Section 2.2, particle filters can be defined on the path space and thus yield a sample $(\bar{x}_{0:t}^k, w_t^k)_{k=1}^{N_x}$ of trajectories approximating the distribution of the paths given the observations and the parameter. In the pseudo-code, “drawing a path” means randomly selecting one of these paths with probabilities proportional to their weights $w_t^{1:N_x}$. Extracting \hat{Z}_t means computing the estimator of Eq. (12). The proposal distribution $q(d\theta^* | \theta)$ can be a Gaussian distribution centered on θ . Intuitively, if the number N_x of particles was infinite, then drawing a path among the path particles would be equivalent to perfect sampling from the filtering distribution $p(dx_{0:t} | y_{0:t}, \theta^*)$, and the likelihood estimator $\hat{Z}_t(\theta^*)$ would yield a perfect evaluation of the likelihood $p(y_{0:t} | \theta^*)$. Thus the algorithm would be a standard Metropolis-Hastings with proposal distribution $q(d\theta^* | \theta)p(dx_{0:t}^* | y_{0:t}, \theta^*)$ and target distribution $\pi_{\theta,t}(d\theta)p(dx_{0:t} | y_{0:t}, \theta)$. The remarkable result of [6] is that for any finite N_x , the PMMH algorithm also generates a Markov chain with invariant distribution $\pi_{\theta,t}(d\theta)p(dx_{0:t} | y_{0:t}, \theta)$. Unsurprisingly larger values of N_x yield better performance of the algorithm and convergence properties of particle MCMC methods have been studied in [3, 4, 7, 27, 64]. In terms of computational cost, the number of particles N_x has to be chosen proportional to t in order to control the variance of the likelihood estimator. Thus each step of PMMH costs t^2 , as was conjectured in [6] and more formally studied in [42, 81]. It is less clear how the number of iterations N_θ must be chosen as a function of

the number of observations t , although some results obtained for standard MCMC could be informative [10]. Under the rather optimistic assumption that N_θ can be chosen independently of t , then the algorithm would overall be of quadratic cost with respect to the number of observations.

Algorithm 2 Particle marginal Metropolis-Hastings.

- 1: Set some $\theta^{(1)}$.
- 2: Run a particle filter with N_x particles given $\theta^{(1)}$.
- 3: Extract $\widehat{Z}_t(\theta^{(1)})$ and draw one path $x_{0:t}^{(1)}$.
- 4: **for** $i = 2$ to N_θ **do**
- 5: Propose $\theta^* \sim q(d\theta^* | \theta^{(i-1)})$.
- 6: Run a particle filter with N_x particles given θ^* .
- 7: Extract $\widehat{Z}_t(\theta^*)$ and draw one path $x_{0:t}^*$.
- 8: Compute:

$$\alpha = \min \left(1, \frac{\widehat{Z}_t(\theta^*) p(\theta^*)}{\widehat{Z}_t(\theta^{(i-1)}) p(\theta^{(i-1)})} \frac{q(\theta^{(i-1)} | \theta^*)}{q(\theta^* | \theta^{(i-1)})} \right).$$

- 9: Set $(\theta^{(i)}, x_{0:t}^{(i)}) = \begin{cases} (\theta^*, x_{0:t}^*) & \text{with probability } \alpha, \\ (\theta^{(i-1)}, x_{0:t}^{(i-1)}) & \text{with probability } 1 - \alpha. \end{cases}$

- 10: **end for**

Various techniques can be used to process the output of MCMC methods to compute estimators of the model evidence as in Eq. (7) [19, 48]. Evidence estimators based on particle MCMC outputs have been proposed in [73]. Thus particle MCMC methods constitute the first class of methods providing practical and consistent approximations of the objects of interest mentioned in Section 1.2 in the context of general implicit models. By design they are iterating over the full dataset $y_{0:t}$, and thus constitute “offline” or “batch” methods, as opposed to the sequential and online features described in Section 1.4. Upon the arrival of a new observation y_{t+1} , the algorithm has to be run again from the beginning. The result of a previous run given $y_{0:t}$ might only be used to design the proposal distribution $q(d\theta^* | \theta)$ and to choose the number of particles N_x .

The SMC² algorithm has been introduced to address this issue [29, 47], and to take one step towards exact, online plug and play methods for Bayesian inference in implicit models. The method processes the observations one after the other, and provides at each step an updated estimator of the various quantities of interest such as the ones described in Section 1.2.

3. A SEQUENTIAL PLUG AND PLAY ALGORITHM

In the light of particle MCMC methods, which mimick the behaviour of ideal MCMC methods when N_x goes to infinity, the idea of SMC² is to mimick an ideal sequential Monte Carlo (SMC) sampler [25, 36] in the setting of hidden Markov models.

3.1. SMC samplers

We first describe the SMC sampler algorithm that we would like to imitate in the hidden Markov model setting. It has been originally proposed for simpler models where it is possible to evaluate the incremental likelihood functions $p(y_t | y_{0:t-1}, \theta)$, for all θ and all t . This is typically the case in parametric models for independent observations, where $p(y_t | y_{0:t-1}, \theta) = p(y_t | \theta)$. Since the full likelihood can be expressed as a product of those incremental likelihoods, it can be evaluated point-wise, and thus the standard Metropolis-Hastings (MH) algorithm is applicable in this context to sample from the posterior distribution $\pi_{\theta,t}(d\theta)$ for a fixed dataset $y_{0:t}$. SMC samplers approximate each posterior distribution $\pi_{\theta,t}(d\theta)$ sequentially over the time t , that is, upon the arrival of new pieces of information. The notion of time can be purely artificial here, e.g. when the data correspond to measurements of different individuals. An adaptive SMC sampler is described in

Algorithm 3. It produces a weighted sample $(\theta_t^k, \omega_t^k)_{k=1}^{N_\theta}$ that approximates the posterior distribution $\pi_{\theta,t}(d\theta)$ at each time t .

Algorithm 3 Adaptive SMC sampler.

- 1: Draw for each $k \in \{1, \dots, N_\theta\}$ $\theta_0^k \sim \pi_\theta(d\theta)$.
 - 2: Set for each k , $\omega_{-1}^k = N_\theta^{-1}$.
 - 3: **for** $t = 0$ to T **do**
 - 4: **if** $\text{ESS}(\omega_{t-1}^{1:N_\theta}) < c$ **then**
 - 5: Construct a proposal distribution q_{t-1} based on the particles $(\theta_{t-1}^k, \omega_{t-1}^k)_{k=1}^{N_\theta}$.
 - 6: [resampling] Sample ancestors $a_t^{1:N_\theta} \sim r(da^{1:N_\theta} | w_{t-1}^{1:N_\theta})$.
 - 7: For each k , replace θ_{t-1}^k by $\theta_{t-1}^{a_t^k}$.
 - 8: Set for each k , $\omega_{t-1}^k = N_\theta^{-1}$.
 - 9: [move] Perform an MCMC move on each particle θ_{t-1}^k using q_{t-1} , leaving $\pi_{\theta,t-1}(d\theta)$ invariant.
 - 10: **end if**
 - 11: Set for each k , $\theta_t^k = \theta_{t-1}^k$.
 - 12: [weighting] Update for each k , $W_t^k = \omega_{t-1}^k p(y_t | y_{0:t-1} \theta_t^k)$.
 - 13: Normalize for each k , $\omega_t^k = W_t^k / \sum_{k=1}^{N_\theta} W_t^k$.
 - 14: **end for**
-

In the algorithm, the resampling step is similar to the one of Algorithm 1, but it is triggered only when the effective sample size (ESS) falls below a threshold c . The ESS is an assessment of the degeneracy of the weights and takes values between 0 and 1. It is defined as the following function of the normalized weights:

$$\text{ESS}(\omega^{1:N_\theta}) = \frac{1}{N_\theta \sum_{k=1}^{N_\theta} (\omega^k)^2}.$$

This adaptive resampling scheme could be applied to the particle filter of Section 2.2 as well, but it proves crucial for SMC samplers, for complexity reasons that will become clear in Section 3.3. The combination of the resampling and the move steps is called the rejuvenation step.

A simple choice of move step leaving $\pi_{\theta,t}(d\theta)$ invariant is to apply a MH kernel with independent proposals from $q_t(d\theta)$. The proposal distribution can be a Gaussian distribution with mean and variance taken as the empirical mean and variance of the particles $(\theta_t^k, \omega_t^k)_{k=1}^{N_\theta}$ [25]. Note that under Bernstein-Von Mises conditions, the posterior distribution converges itself to a Gaussian distribution, and thus using an adaptive Gaussian proposal distribution in the rejuvenation step is an asymptotically optimal choice. The move step then consists in applying one step of MH to each of the N_θ particles.

The algorithmic parameters left to choose are the number of particles N_θ and the ESS threshold c , which can be set to 50% by default. Higher values mean more rejuvenation steps, which constitute the bulk of the computational cost of the algorithm. The algorithm has been extended to a more general form in [36], which allows various algorithmic improvements as well as a unified theoretical study under the Feynman-Kac framework. Various articles study its theoretical properties [57, 80, 88]. In particular [79] demonstrates its theoretical advantage over MCMC when the posterior distribution is multimodal. The effect of triggering resampling steps based on an ESS criterion has been studied in [37]. The behaviour of the algorithm with respect to the dimension d_θ of the parameter space has been studied in [14].

Similarly to the likelihood estimator given by particle filters in Eq. (12), SMC samplers yield an estimator $\widehat{\mathcal{Z}}_t$ of the evidence \mathcal{Z}_t defined in Eq. (7), that can be computed as

$$\widehat{\mathcal{Z}}_t = \prod_{s=0}^t \left(\sum_{k=1}^{N_\theta} \omega_{s-1}^k p(y_s | y_{0:s-1}, \theta_s^k) \right). \quad (13)$$

The form of the estimator is slightly different from the likelihood estimator in Eq. (12) because the resampling steps are not applied at every step. One way to justify it is to consider that if $(\theta_t^k, \omega_{t-1}^k)_{k=1}^{N_\theta}$ is a consistent particle approximation of $\pi_{\theta, t-1}(d\theta)$, e.g. in probability, and remembering that the weights $\omega_{t-1}^{1:N_\theta}$ are normalized in the algorithm, then

$$\sum_{k=1}^{N_\theta} \omega_{t-1}^k p(y_t | y_{0:t-1}, \theta_t^k) \xrightarrow[N_\theta \rightarrow \infty]{\mathbb{P}} \int_{\Theta} p(y_t | y_{0:t-1}, \theta) \pi_{\theta, t-1}(d\theta) = p(y_t | y_{0:t-1}).$$

Taking the product over time steps yields an estimator of the full evidence $p(y_{0:t})$, given the model. The inclusion of SMC samplers into the Feynman-Kac framework of general particle methods allows the study of the properties of this estimator, in particular [36] obtain a central limit theorem. Empirically [94] demonstrate the advantage of SMC samplers over MCMC methods to estimate the model evidence.

3.2. An approximate SMC sampler for hidden Markov models

The original SMC sampler as in Algorithm 3 cannot be directly applied to the hidden Markov model scenario, in the same way that standard MH could not be applied: the likelihood function as defined in Eq. (4) cannot be evaluated point-wise, and the incremental likelihoods $p(y_t | y_{0:t-1}, \theta)$ cannot either. Mimicking the reasoning behind particle MCMC methods, particle filters can be used to obtain estimators of those likelihood terms. For simplicity, we present an SMC sampler algorithm to sample from $\pi_{\theta, t}(d\theta)$ only, but the same algorithm can be used to sample from the joint distribution $\pi_{\theta, t}(d\theta)p(dx_{0:t} | y_{0:t}, \theta)$, as shown in [29]. The article [47] essentially proposed the same algorithm independently, with challenging applications in econometrics.

To each of the N_θ parameter values produced by the SMC sampler as in Algorithm 3, we thus attach a particle filter with N_x particles as in Algorithm 1, hence the name SMC² evoking those two layers of particle approximations. To avoid confusion we will talk about θ -particles and x -particles, and denote respectively by N_θ and N_x their numbers. At any time t , each of the N_θ θ -particles is indexed by k as in θ_t^k , while each of the associated x -particles is indexed by n, k as in $x_t^{n, k}$. The method is described in Algorithm 4.

The algorithm follows the structure of Algorithm 3, except that each θ -particle is equipped with a particle filter that is updated at each step. The differences are summarised in the following two points.

- At time t , the weight of the θ -particle θ_t^k is updated using an estimator $\hat{p}(y_t | y_{0:t-1}, \theta_t^k)$ obtained from the associated particle filter, instead of the true incremental likelihood $p(y_t | y_{0:t-1}, \theta_t^k)$.
- The move step to rejuvenate the θ -particles relies on particle MCMC instead of MCMC.

A more complete description of the algorithm is given in [29]. Let us simply mention that SMC² is a standard SMC sampler targeting an extended distribution which admits $\pi_{\theta, t}(d\theta)p(dx_{0:t} | y_{0:t}, \theta)$ as one of its marginal distributions, for any number N_x . Hence, the algorithm falls into the class of exact approximations, similarly to particle MCMC methods. Thus filtering under parameter uncertainty as defined in Eq. (6) can be addressed consistently, for a finite N_x and N_θ going to infinity, and furthermore the algorithm is sequential by design. Before turning to its computational complexity, let us mention that the model evidence can be retrieved with the following estimator

$$\widehat{\mathcal{Z}}_t = \prod_{s=0}^t \left(\sum_{k=1}^{N_\theta} \omega_{s-1}^k \hat{p}(y_s | y_{0:s-1}, \theta_s^k) \right). \quad (14)$$

Thus the algorithm can consistently compute integrals such as Eq. (9).

3.3. Complexity of SMC samplers

Going back to the ideal SMC sampler in Algorithm 3, a MH move for each particle θ_t^k at time t requires an evaluation of the likelihood $p(y_{0:t} | \theta_t^k)$, which costs $\mathcal{O}(t)$. If a rejuvenation step was performed at every step from time 0 to t , the algorithm would then cost $\mathcal{O}(N_\theta t^2)$. Fortunately, the ESS decreases slower and slower, and thus the rejuvenation step occurs less and less often, hence the cost typically reduces to $\mathcal{O}(N_\theta t)$, as shown

Algorithm 4 SMC² sampler.

-
- 1: [θ -initialization] Draw for each $k \in \{1, \dots, N_\theta\}$ $\theta_0^k \sim \pi_\theta(d\theta)$.
 - 2: [x -initialization] For each k , draw for each $n \in \{1, \dots, N_x\}$ $x_0^{n,k} \sim \mu(dx_0 | \theta_0^k)$.
 - 3: Set for each k , $\omega_{-1}^k = N_\theta^{-1}$.
 - 4: **for** $t = 0$ to T **do**
 - 5: **if** $\text{ESS}(\omega_{t-1}^{1:N_\theta}) < c$ **then**
 - 6: Construct a proposal distribution q_{t-1} based on the θ -particles $(\theta_{t-1}^k, \omega_{t-1}^k)_{k=1}^{N_\theta}$.
 - 7: [θ -resampling] Sample ancestors $a_t^{1:N_\theta} \sim r(da^{1:N_\theta} | w_{t-1}^{1:N_\theta})$.
 - 8: For each k , replace θ_{t-1}^k by $\theta_{t-1}^{a_t^k}$.
 - 9: Set for each k , $\omega_{t-1}^k = N_\theta^{-1}$.
 - 10: [θ -move] Perform a PMMH move on each θ -particle θ_{t-1}^k using q_{t-1} , leaving $\pi_{\theta,t-1}(d\theta)$ invariant.
 - 11: **end if**
 - 12: Set for each k , $\theta_t^k = \theta_{t-1}^k$.
 - 13: [x -weighting] Compute for each k and n , $w_t^{n,k} = g(y_t | x_t^{n,k}, \theta_t^k)$.
 - 14: [x -resampling] Sample for each k , $a_t^{1:N_x,k} \sim r(da^{1:N_x} | w_t^{1:N_x,k})$.
 - 15: [x -transition] Draw for each k and n , $x_{t+1}^{n,k} \sim f(dx_{t+1} | x_t^{a_t^{n,k}}, \theta_t^k)$.
 - 16: Compute for each k , $\hat{p}(y_t | y_{0:t-1}, \theta_t^k) = N_x^{-1} \sum_{n=1}^{N_x} w_t^{n,k}$.
 - 17: [θ -weighting] Update for each k , $W_t^k = \omega_{t-1}^k \hat{p}(y_t | y_{0:t-1}, \theta_t^k)$.
 - 18: Normalize for each k , $\omega_t^k = W_t^k / \sum_{k=1}^{N_\theta} W_t^k$.
 - 19: **end for**
-

in Theorem 1 of [25]. In other words, at each step t , either the assimilation of y_t costs $\mathcal{O}(t)$ or $\mathcal{O}(1)$, whether or not a rejuvenation step is performed, which happens with a probability decreasing with t . Let us denote by p_t the probability of a rejuvenation step occurring at time t . If the other operations are of cost 1 at each step, then the overall cost C_t for the algorithm to reach step t satisfies

$$\mathbb{E}[C_t] = \sum_{s=0}^t (p_s \times s + (1 - p_s) \times 1) = (t+1) + \sum_{s=0}^t p_s \times (s-1)$$

which indeed is linear in t if $p_t = \mathcal{O}(1/t)$. This is another formulation of the result in [25]. Thus the algorithm is online in the sense described in Section 1.4. Note that the algorithm requires more and more memory, as a rejuvenation step at time t involves browsing over the past dataset $y_{0:t}$, which thus has to be kept available. Thus the algorithm is not “online” memory-wise but only in terms of computational cost. One can hope that the errors are uniformly bounded over time for a fixed N_θ if the rejuvenation steps are performing equally well across all time steps. This motivates the adaptation of the proposal distribution q_t in Algorithm 3.

For the SMC² algorithm of Algorithm 4, the same reasoning applies, motivated by empirical results such as Figure 4 to be described in the next section. The difference is that, in order to bound the errors uniformly over time, and for the particle MCMC steps to perform equally well across all time steps, one needs to increase the number N_x of x -particles with t . Scaling N_x linearly with t , the cost of running N_θ particle filters with N_x x -particles for t steps is $\mathcal{O}(N_\theta t^2)$. Under the same occurrence pattern of rejuvenation steps, [29] obtain an overall computational cost in $\mathcal{O}(N_\theta t^2)$. The SMC² algorithm is thus sequential but not online. Upon the arrival of a new piece of observation, the estimator can be updated, but one has to increase the computational effort linearly with the number of observations in order to obtain time uniform guarantees. A modification of SMC² is proposed in [29] so that N_x can be automatically increased along the observations when required. The acceptance rate of the rejuvenation steps are monitored to assess whether N_x is large enough at any time t . This modification does not make the algorithm online, but allows the automatic adjustment of the computational cost to guarantee a stable performance over time.

There currently exists no method to perform online and exact Bayesian inference for general hidden Markov models [30, 93], which poses a serious challenge in the presence of very long time series. In terms of scaling with the number of particles, SMC algorithms are very amenable to modern parallel architectures. The algorithm is typically of linear complexity in N_x and in N_θ , and most of the computation can be done in parallel, except for the resampling steps. This has motivated a series of articles in the recent years, both in the computational literature [16, 53, 70] and in the methodological literature [23, 87, 90].

4. NUMERICAL ILLUSTRATION

The PZ model of Section 1.3 is used to illustrate the various outputs of SMC². Given the parameters set in Section 1.3, $T = 365$ observations are generated as shown on Figure 2. The algorithm is run with $N_\theta = 1024$, $N_x = 1024$ and an ESS threshold c of 50%. The proposal distribution q_t of the move steps is a Gaussian distribution using the empirical mean and covariance of the weighted particles $(\theta_t^k, \omega_t^k)_{k=1}^{N_\theta}$. Each rejuvenation step performs five successive PMMH moves on each particle θ_t^k . The resampling distribution for both the θ -particles and the x -particles is chosen to be the systematic resampling scheme [21]. To diagnose the behaviour of an SMC² run, the ESS of the θ -particles is plotted against time on Figure 3. As expected the ESS decreases slower and slower with the time steps, resulting in only three rejuvenation steps in the second half of the dataset, whereas ten rejuvenation steps occurred in the first half. The acceptance rate of the move steps is found to be above 40% at the end of the run.

The computational cost of the algorithm is represented on Figure 4. More precisely what is plotted is the number of times that the transition f and the measurement g are called, for each of the 1024 θ -particles. Since there are 365 time steps and 1024 x -particles per θ -particle, if no PMMH move was performed there would be $365 \times 1024 \approx 3.7 \times 10^5$ transitions per θ -particle. Since five PMMH steps are performed at each rejuvenation step, the number of calls per θ -particle reaches 6×10^6 . The dashed line represents a linear regression of the number of calls against time, indicating that the number of calls seems to grow linearly in t . Note that this linear trend is obtained for a fixed $N_x = 1024$. The quadratic cost of the overall method mentioned in Section 3.3 comes from the fact that one would eventually need to increase N_x if observations kept arriving. Thus if we had two years of daily data instead of one, and if we wanted to obtain the same relative error in estimating the integrals of Section 1.2, we would set N_x to 2048 and the overall expected computational time would be multiplied by four. Since there are 1024 θ -particles, the total number of calls to the functions f and g is in the billions. For the PZ model, each transition involves solving numerically a differential equation, here using a Runge-Kutta method RK4(3)5[2R+] as in [69]. In wall-clock time, this SMC² run took about 50 minutes on a standard desktop computer with 8 cores using the optimized software LibBi [68]. Across runs, the random occurrence of rejuvenation steps incurs random computing times. We collected runtimes between 40 and 60 minutes, using the same algorithmic parameters, over five independent runs.

The approximation of the posterior distribution $\pi_{\theta,T}(d\theta)$ at the final time $T = 365$ is represented by the pairwise contour plots of Figure 5. We see that the posterior distribution concentrates in the neighborhood of the parameter used to generate the dataset, indicated by black dots on the figure. We note negative correlations between some of the parameters, in particular between m_l and m_q , which both explain the instantaneous decrease of the zooplankton population size, and between σ_y and σ_α , which both account for the stochasticity of the model. We observe that the mode of the posterior distribution is not exactly located at the data-generating parameter because the inference is conditional upon a finite number of observations. Indeed, one could only expect the data-generating parameter to be recovered when the number of observations goes to infinity. Note also that uniform priors have been used for all the parameters, therefore the mode of the posterior distribution corresponds exactly to the maximum likelihood estimate.

One advantage of sequential inference is the ability to investigate each intermediate posterior distribution $\pi_{\theta,t}(d\theta)$ for $t = 0, \dots, T$. Figure 6 represents this evolution for the first 50 time steps and for each parameter. The grey ribbons indicate the 10%, 20%, 30%, 40%, 60%, 70%, 80%, 90% quantiles of each marginal posterior distribution. The dashed lines indicate the values used to generate the dataset. We see the posterior distributions going nearer the data-generating parameter as more observations are being assimilated. We also observe that




FIGURE 3. Effective Sample Size against time, over one run of SMC^2 on the PZ model. The vertical dashed lines represent the resampling times and the horizontal dashed line represents the ESS threshold, set to 50% of N_θ .




FIGURE 4. Cumulative cost per θ -particle during one run of SMC^2 on the PZ model, with $N_x = 1024$ and five PMMH moves per rejuvenation step. The cost is measured in the number of calls to the transition sampling function and the measurement density function. The dashed line represents a linear regression of the cost over the time index.

this concentration occurs at a different rate for each parameter. Indeed, according to asymptotic results on the posterior distribution (Chapter 1 of [49]), the asymptotic concentration rates depend on the Fisher information matrix of the model. In a non-asymptotic setting, as is the case in practice, we could imagine using plots similar to Figure 6 to guess how many more observations would be needed to reach a given precision for each parameter.

The possibility to perform prediction under parameter uncertainty is illustrated on Figure 7. At every time step, an 80% predictive region is inferred from the particle approximation of y_{t+1} given $y_{0:t}$. The successive regions are joined together in a grey ribbon. The actual observations are plotted as circles if they fall in the predictive region, and triangles if they fall outside. At the end of the run, 77 observations have landed outside the predictive region, which represents 21% instead of the targeted 20%. Since the observations are generated from the model, it is expected that asymptotically in t , 20% would fall outside the 80% predictive region. Figure 8 is a close-up of Figure 7, focusing on the first 50 time steps.




FIGURE 5. Posterior distribution of the parameters of the PZ model given the synthetic dataset of 365 observations. Contour lines represent the estimated density of the pairwise marginal distributions $(\sigma_\alpha, \mu_\alpha)$, $(\sigma_\alpha, \sigma_y)$ and (m_l, m_q) . The dots indicate the values used to generate the dataset.

The model evidence estimator of Eq. (14) can be illustrated by introducing another model. We consider a simplified model PZ^* , which is defined as model PZ except that the quadratic mortality term is removed from the differential equation:

$$\begin{aligned}\frac{dp_t}{dt} &= \alpha p_t - cp_t z_t, \\ \frac{dz_t}{dt} &= ecp_t z_t - m_l z_t.\end{aligned}$$

Thus the parameter is $\theta = (\mu_\alpha, \sigma_\alpha, \sigma_y, m_l)$ and we use the same uniform prior distributions as for the PZ model. We put a uniform prior over the two models and thus the posterior odds as in Eq. (10) reduce to the Bayes factor, $p(y_{0:t} | PZ)/p(y_{0:t} | PZ^*)$. This ratio can be obtained by approximating the evidence using the estimator of Eq. (14) for each model. The same algorithmic parameters as described above are used for each model.

Figure 9 shows the Bayes factors against time, obtained from five independent runs, and Figure 10 is a close-up on the first 100 time steps. The bottom horizontal dashed line indicates 1. A Bayes factor of 1 indicates no support of the data for one model compared to the other. Values close to zero support model PZ^* while values larger than one support model PZ. The graph shows that for the first 50 observations, model PZ^* seems supported by the data. However with more observations, the Bayes factor starts supporting the PZ model, and after time 100 each of the five independent runs estimates the factor above 100. The factor keeps increasing to extremely large values as more observations are assimilated. Here the dataset is generated using the PZ model so the end result does not come as a surprise. The sequential Bayes factor estimation shows Occam's razor principle in action, as mentioned in Section 1.2: the simpler model is favoured when few observations are available. Here about 100 data points are enough to choose the data-generating model with confidence, according to the Bayes factor criterion. Note how the five independent evidence estimates diverge from each other as observations accrue, showing that the estimator is not stable over time for a fixed number of particles N_x and N_θ . This confirms that the evidence estimator from SMC² is not online.

5. DISCUSSION

Let us discuss again the objects of inference of Section 1.2 in the light of the reviewed methodology. For a given parameter value, filtering and prediction as in the integrals of Eq. (1), (2), (3) can be approximated in an

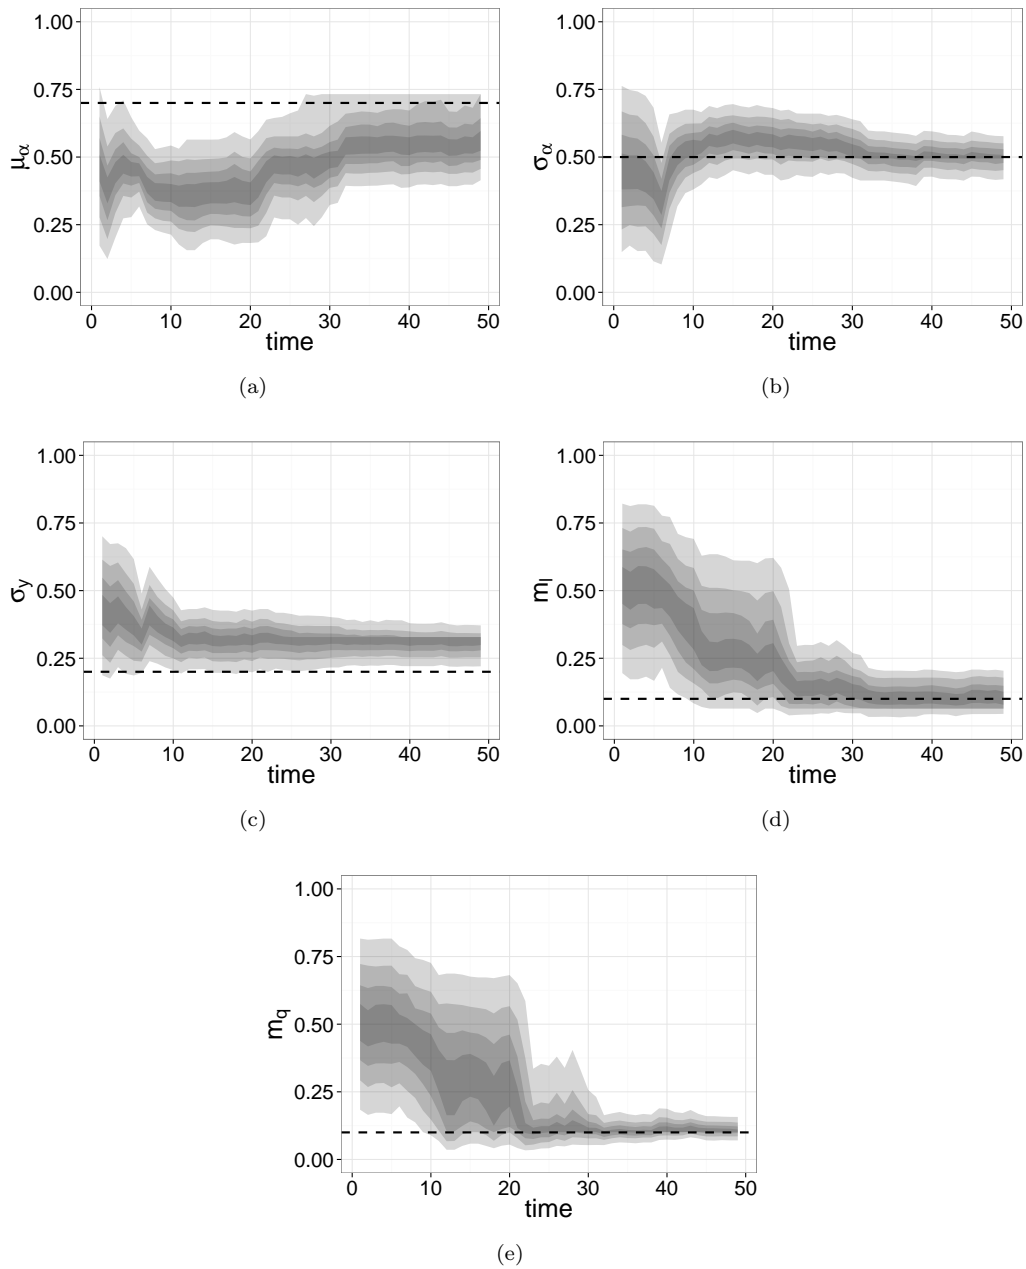


FIGURE 6. Evolution of the posterior distribution of each parameter of the PZ model obtained by SMC², over the first 50 time steps. The grey ribbons indicate the 10%, 20%, 30%, 40%, 60%, 70%, 80%, 90% quantiles. The dashed lines indicate the values used to generate the dataset.




FIGURE 7. One step predictions obtained using one run of SMC^2 on the PZ model. The dark ribbon indicates the 80% predictive region of y_{t+1} given $y_{0:t}$ for each time, under parameter uncertainty. Observations that land in the predictive region are indicated by circles, whereas observations landing outside are indicated by triangles.




FIGURE 8. Same as Figure 7, but limited to the first 50 time steps. The grey ribbons indicate the quantiles of one-step predictive regions.

online manner using particle filters, as described in Section 2. There are currently no online and exact methods, in the sense of Section 1.4, that take into account parameter uncertainty as in Eq. (5), (6), (7). The SMC^2 method proposed in [29] and [47], and described in Section 3, is sequential, in the sense that the estimators can be updated upon the arrival of new observations. However the incremental cost of the algorithm has to grow linearly with t in order to control the relative variance of the estimators. Hence a complete run of the algorithm has a quadratic cost in the length of the time series, and thus is not applicable for long time series. Informally, for the PZ model used as an illustration in Section 4, the SMC^2 algorithm runs in a reasonable time on standard hardware, for thousands of observations, but not for millions. The numbers would change with the model and the application, but in general online inference under parameter uncertainty is still an open question and an active area of research; see [30, 93] for recent developments.

One of the difficulties comes from the likelihood estimator of Eq. (12), which requires a quadratic cost in the number of observations to guarantee a bounded relative error. On the other hand, the Kalman filter yields likelihood evaluations in a linear cost, but only for linear Gaussian models. It is unclear whether an intermediate




FIGURE 9. Bayes factor of the PZ model versus the simplified PZ* model against time. The bottom dashed line indicates 1 and the top one indicates 100. Values larger than 1 indicate support for the PZ model. The full lines correspond to the estimates of the Bayes factor for five independent runs of the SMC² algorithm.




FIGURE 10. Same as Figure 9, but limited to the first 100 time steps. Initially the simpler PZ* model is preferred by the Bayes factor, but after 100 observations the more complex, data-generating PZ model becomes strongly supported by the criterion.

setting exists, where the likelihood could be estimated in a super-linear but sub-quadratic cost, at least for some models.

For time series of moderate length, filtering, prediction and parameter inference are still challenging problems when the dimension d_x of the state space \mathbb{X} is large. Indeed the variance of standard particle filter estimators typically increases exponentially with d_x . Recent developments such as [15, 74] could help scaling particle methods to larger dimensions, with many potential applications in spatial state space models [24]. Another issue specific to large-dimensional state space models is the large computer memory required to store the particles, and especially the paths $(\bar{x}_{0:t}^{1:N_x}$ in the notation of Section 2.2). Large memory usage also involves large communication costs on distributed hardware, whenever particles have to be sent from one machine to another. The expected memory usage of storing the paths has been studied in [56]. Methods to reduce the memory and communication costs on distributed hardware have been proposed in [77, 78].

Another challenge is to adapt computational methods to larger classes of models. The plug and play methods described in Section 2 are compatible with parametric models, where the latent process can be simulated and

the measurement density can be evaluated point-wise. On top of hidden Markov models, particle filters can be implemented for non-Markovian models, as long as these two requirements are met. Inference in non-Markovian models using particle methods has been recently considered in [63]. The performance of particle methods in non-Markovian settings has been partially studied in [28]. Recent applications of non-Markovian models, for instance in probabilistic programming [91], motivate further research in this direction.

A number of articles have considered non-parametric hidden Markov models. The authors of [20] consider linear models with a transition equation of the form $x_t = A_t x_{t-1} + G_t v_t$, where the distribution of the noise v_t is modelled with a Dirichlet process mixture. A non-parametric model is also considered for the measurement noise. Estimation is then performed using Markov chain Monte Carlo or sequential Monte Carlo methods. In the more recent literature, [46] consider a transition equation of the form $x_t = f(x_{t-1}) + v_t$ and put a Gaussian process prior on the function f ; particle Markov chain Monte Carlo methods then enable inference under parameter uncertainty. Note that combining non-parametric models for the function f and for the noise v_t is not obvious because of identifiability issues. Other instances of non-parametric hidden Markov models consider the case where the hidden process lives on an infinite but discrete state space [44, 84]. Particle Markov chain Monte Carlo methods have recently been used in this context [86]. The case of finite state space and non-parametric specification of the measurement distribution is considered in [92]. Sequential algorithms to perform inference in continuous space, non-linear, non-parametric hidden Markov models would constitute an interesting addition to the current methodology.

The author gratefully acknowledges EPSRC for funding this research through grant EP/K009362/1, thanks the organizers of the Journées MAS 2014, thanks Arnaud Doucet, Lawrence Murray and Aimee Taylor for useful comments. This article is dedicated to the memory of Philip Perry and some long discussions on the Bayesian approach.

REFERENCES

- [1] J.A. Achcar, E.R. Rodrigues, and G. Tzintzun. Using stochastic volatility models to analyse weekly ozone averages in Mexico City. *Environmental and Ecological Statistics*, 18(2):271–290, 2011.
- [2] B.D.O. Anderson and J.B. Moore. *Optimal filtering*. Courier Dover Publications, 2012.
- [3] C. Andrieu and M. Vihola. Establishing some order amongst exact approximations of MCMCs. *arXiv preprint arXiv:1404.6909*, 2014.
- [4] C. Andrieu and M. Vihola. Convergence properties of pseudo-marginal Markov chain Monte Carlo algorithms. *The Annals of Applied Probability*, 25(2):1030–1077, 2015. URL <http://dx.doi.org/10.1214/14-AAP1022>.
- [5] C. Andrieu, A. Doucet, and V. Tadic. On-line parameter estimation in general state-space models. In *Proceedings of the 44th Conference on Decision and Control*, pages 332–337, 2005.
- [6] C. Andrieu, A. Doucet, and R. Holenstein. Particle Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society: Series B*, 72(3):269–342, 2010.
- [7] C. Andrieu, A. Lee, and M. Vihola. Uniform ergodicity of the iterated conditional SMC and geometric ergodicity of particle Gibbs samplers. *arXiv preprint arXiv:1312.6432*, 2013.
- [8] O.E. Barndorff-Nielsen and N. Shephard. Econometric analysis of realized volatility and its use in estimating stochastic volatility models. *Journal of the Royal Statistical Society: Series B*, 64(2):253–280, 2002.
- [9] M.A. Beaumont. Approximate Bayesian computation in evolution and ecology. *Annual review of ecology, evolution, and systematics*, 41:379–406, 2010.
- [10] A. Belloni and V. Chernozhukov. On the computational complexity of MCMC-based estimators in large samples. *The Annals of Statistics*, pages 2011–2055, 2009.
- [11] J.M. Bernardo and A.F.M. Smith. *Bayesian theory*, volume 405. John Wiley & Sons, 2009.
- [12] A.A. Berryman. The origins and evolution of predator-prey theory. *Ecology*, 73(5):1530–1535, 1992. URL <http://www.jstor.org/stable/1940005>.
- [13] F. Bertoli and A.N. Bishop. Adaptively blocked particle filtering with spatial smoothing in large-scale dynamic random fields. *arXiv preprint arXiv:1406.0136*, 2014.

- [14] A. Beskos, D. Crisan, and A. Jasra. On the stability of sequential Monte Carlo methods in high dimensions. *The Annals of Applied Probability*, 24(4):1396–1445, 2014. URL <http://dx.doi.org/10.1214/13-AAP951>.
- [15] A. Beskos, D. Crisan, A. Jasra, K. Kamatani, and Y. Zhou. A stable particle filter in high-dimensions. *arXiv preprint arXiv:1412.3501*, 2014.
- [16] M. Bolic, P.M. Djuric, and S. Hong. Resampling algorithms and architectures for distributed particle filters. *IEEE Transactions on Signal Processing*, 53(7):2442–2450, 2005.
- [17] C. Bretó, D. He, E.L. Ionides, and A.A. King. Time series analysis via mechanistic models. *The Annals of Applied Statistics*, 3(1):319–348, 2009.
- [18] O. Cappé, E. Moulines, and T. Rydén. *Inference in Hidden Markov Models*. Springer-Verlag, New York, 2005.
- [19] B.P. Carlin and S. Chib. Bayesian model choice via Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society: Series B*, pages 473–484, 1995.
- [20] F. Caron, M. Davy, A. Doucet, E. Duflos, and P. Vanheeghe. Bayesian inference for linear dynamic models with Dirichlet process mixtures. *IEEE Transactions on Signal Processing*, 56(1):71–84, 2008.
- [21] J. Carpenter, P. Clifford, and P. Fearnhead. Improved particle filter for non-linear problems. *IEE Proceedings Radar, Sonar and Navigation*, 146(1):2–7, 1999.
- [22] F. Cérou, P. Del Moral, and A. Guyader. A non-asymptotic theorem for unnormalized Feynman–Kac particle models. *Annales de l’Institut Henri Poincaré*, 47(3):629–649, 2011.
- [23] H.P. Chan, C.W. Heng, and A. Jasra. Theory of parallel particle filters for hidden Markov models. *arXiv preprint arXiv:1409.4160*, 2014.
- [24] Y. Chen and S. Reich. Data assimilation: a dynamical system perspective. *Frontiers in Applied Dynamical Systems*, to appear, 2015.
- [25] N. Chopin. A sequential particle filter for static models. *Biometrika*, 89:539–552, 2002.
- [26] N. Chopin. Central Limit Theorem for sequential Monte Carlo methods and its application to Bayesian inference. *The Annals of Statistics*, 32(6):2385–2411, 2004.
- [27] N. Chopin and S.S. Singh. On the particle Gibbs sampler. *Bernoulli*, to appear, 2015.
- [28] N. Chopin, P. Del Moral, and S. Rubenthaler. Stability of Feynman–Kac formulae with path-dependent potentials. *Stochastic Processes and their Applications*, 121(1):38–60, 2011.
- [29] N. Chopin, P.E. Jacob, and O. Papaspiliopoulos. SMC²: an efficient algorithm for sequential analysis of state space models. *Journal of the Royal Statistical Society: Series B*, 75(3):397–426, 2013.
- [30] D. Crisan and J. Miguez. Nested particle filters for online parameter estimation in discrete-time state-space Markov models. *arXiv preprint arXiv:1308.1883*, 2013.
- [31] A.P. Dawid. Statistical theory: the prequential approach. *Journal of the Royal Statistical Society: Series A*, pages 278–292, 1984.
- [32] T.A. Dean, S.S. Singh, A. Jasra, and G.W. Peters. Parameter estimation for hidden Markov models with intractable likelihoods. *Scandinavian Journal of Statistics*, 2014.
- [33] P. Del Moral. *Feynman–Kac Formulae*. Springer, 2004.
- [34] P. Del Moral and A. Doucet. On a class of genealogical and interacting Metropolis models. *Séminaire de Probabilités XXXVII*, pages 415–446, 2003.
- [35] P. Del Moral and A. Doucet. Particle methods: an introduction with applications. In *ESAIM: Proceedings*, volume 44, pages 1–46. EDP Sciences, 2014.
- [36] P. Del Moral, A. Doucet, and A. Jasra. Sequential Monte Carlo samplers. *Journal of the Royal Statistical Society: Series B*, 68(3):411–436, 2006.
- [37] P. Del Moral, A. Doucet, and A. Jasra. On adaptive resampling strategies for sequential Monte Carlo methods. *Bernoulli*, 18(1):252–278, 2012. URL <http://dx.doi.org/10.3150/10-BEJ335>.
- [38] R. Douc, A. Garivier, E. Moulines, and J. Olsson. Sequential Monte Carlo smoothing for general state space hidden Markov models. *The Annals of Applied Probability*, 21(6):2109–2145, 2011. URL <http://dx.doi.org/10.1214/10-AAP735>.

- [39] A. Doucet and A.M. Johansen. A tutorial on particle filtering and smoothing: fifteen years later. In *Handbook of Nonlinear Filtering*. Oxford, UK: Oxford University Press, 2011.
- [40] A. Doucet, N. de Freitas, and N.J. Gordon. *Sequential Monte Carlo Methods in Practice*. Springer-Verlag, New York, 2001.
- [41] A. Doucet, G. Poyiadjis, and S.S. Singh. Sequential Monte Carlo computation of the score and observed information matrix in state-space models with application to parameter estimation. *Biometrika*, 98:65–80, 2011.
- [42] A. Doucet, G. Deligiannidis, M.K. Pitt, and R. Kohn. Efficient implementation of Markov chain Monte Carlo when using an unbiased likelihood estimator. *Biometrika*, to appear, 2015.
- [43] P. Fearnhead. Markov chain Monte Carlo, sufficient statistics, and particle filters. *Journal of Computational and Graphical Statistics*, 11(4):848–862, 2002.
- [44] E.B. Fox, E.B. Sudderth, M.I. Jordan, and A.S. Willsky. An HDP-HMM for systems with state persistence. In *Proceedings of the 25th international conference on Machine learning*, pages 312–319. ACM, 2008.
- [45] M. Frei and H.R. Künsch. Bridging the ensemble Kalman and particle filters. *Biometrika*, 100(4):781–800, 2013.
- [46] R. Frigola, F. Lindsten, T.B. Schön, and C. Rasmussen. Bayesian inference and learning in Gaussian process state-space models with particle MCMC. In *Advances in Neural Information Processing Systems*, pages 3156–3164, 2013.
- [47] A. Fulop and J. Li. Efficient learning via simulation: A marginalized resample-move approach. *Journal of Econometrics*, 176(2):146–161, 2013.
- [48] A. Gelman and X.L. Meng. Simulating normalizing constants: From importance sampling to bridge sampling to path sampling. *Statistical Science*, 13(2):163–185, 1998.
- [49] J.K. Ghosh and R.V. Ramamoorthi. *Bayesian nonparametrics*. Springer-Verlag, New York, 2003.
- [50] W.R. Gilks and C. Berzuini. Following a moving target - Monte Carlo inference for dynamic Bayesian models. *Journal of the Royal Statistical Society: Series B*, 63:127–146, 2001.
- [51] N.J. Gordon, D.J. Salmond, and A.F.M. Smith. Novel approach to non-linear/non-Gaussian Bayesian state estimation. *IEE Proceedings F*, 140(2):107–113, 1993.
- [52] J.A. Hoeting, D. Madigan, A.E. Raftery, and C.T. Volinsky. Bayesian model averaging: a tutorial. *Statistical science*, pages 382–401, 1999.
- [53] S. Hong and P.M. Djuric. High-throughput scalable parallel resampling mechanism for effective redistribution of particles. *IEEE Transactions on Signal Processing*, 54(3):1144–1155, 2006.
- [54] E.L. Ionides, C. Bretó, and A.A. King. Inference for non-linear dynamical systems. *Proceedings of the National Academy of Sciences*, 103(49):18438–18443, 2006.
- [55] E.L. Ionides, A. Bhadra, Y. Atchadé, and A.A. King. Iterated filtering. *The Annals of Statistics*, 39(3):1776–1802, 2011.
- [56] P.E. Jacob, L.M. Murray, and S. Rubenthaler. Path storage in the particle filter. *Statistics and Computing*, 25(2):487–496, 2013.
- [57] A. Jasra and A. Doucet. Stability of sequential Monte Carlo samplers via the Foster–Lyapunov condition. *Statistics & Probability Letters*, 78(17):3062–3069, 2008.
- [58] E.M. Jones, J. Parslow, and L.M. Murray. A Bayesian approach to state and parameter estimation in a phytoplankton-zooplankton model. *Australian Meteorological and Oceanographic Journal*, 59:7–16, 2010.
- [59] N. Kantas, A. Doucet, S.S. Singh, and J.M. Maciejowski. An overview of sequential Monte Carlo methods for parameter estimation in general state-space models. In *15th IFAC Symposium on System Identification (SYSID)*, Saint-Malo, France, volume 102, page 117, 2009.
- [60] N. Kantas, A. Doucet, S.S. Singh, J.M. Maciejowski, and N. Chopin. On particle methods for parameter estimation in state-space models. *Statistical science*, to appear, 2015.
- [61] R.E. Kass and A.E. Raftery. Bayes factors. *Journal of the American Statistical Association*, 90(430):773–795, 1995.

- [62] G. Kitagawa. A self-organizing state-space model. *Journal of the American Statistical Association*, 93: 1203–1215, 1998.
- [63] F. Lindsten, T.B. Schön, and M.I. Jordan. Ancestor sampling for particle Gibbs. In *Advances in Neural Information Processing Systems*, pages 2591–2599, 2012.
- [64] F. Lindsten, R. Douc, and E. Moulines. Uniform ergodicity of the particle Gibbs sampler. *Scandinavian Journal of Statistics*, to appear, 2015.
- [65] J. Liu and M. West. Combined parameter and state estimation in simulation-based filtering. In A. Doucet, N. de Freitas, and N. J. Gordon, editors, *Sequential Monte Carlo Methods in Practice*, pages 197–223. Springer-Verlag, 2001.
- [66] D.J.C. MacKay. Bayesian interpolation. *Neural computation*, 4(3):415–447, 1992.
- [67] J.-M. Marin, P. Pudlo, C.P. Robert, and R.J. Ryder. Approximate Bayesian computational methods. *Statistics and Computing*, 22(6):1167–1180, 2012.
- [68] L.M. Murray. Bayesian state-space modelling on high-performance hardware using LibBi. *arXiv preprint arXiv:1306.3277*, 2013.
- [69] L.M. Murray, E.M. Jones, and J. Parslow. On disturbance state-space models and the particle marginal Metropolis–Hastings sampler. *SIAM/ASA Journal of Uncertainty Quantification*, 1(1):484–521, 2013.
- [70] L.M. Murray, A. Lee, and P.E. Jacob. Parallel resampling in the particle filter. *arXiv preprint arXiv:1301.4019*, 2014.
- [71] M. Papaspiliopoulos, O. and Ruggiero. Optimal filtering and the dual process. *Bernoulli*, 20(4):1999–2019, 2014. URL <http://dx.doi.org/10.3150/13-BEJ548>.
- [72] O. Papaspiliopoulos, G.O. Roberts, and M. Sköld. A general framework for the parametrization of hierarchical models. *Statistical Science*, 22(1):59–73, 2007.
- [73] G.W. Peters, G.R. Hosack, and K.R. Hayes. Ecological non-linear state space model selection via adaptive particle Markov chain Monte Carlo (adpMCMC). *arXiv preprint arXiv:1005.2238*, 2010.
- [74] P. Rebeschini and R. Van Handel. Can local particle filters beat the curse of dimensionality? *The Annals of Applied Probability*, to appear, 2015.
- [75] C.P. Robert. *The Bayesian choice: from decision-theoretic foundations to computational implementation*. Springer Verlag, 2007.
- [76] C.P. Robert and G. Casella. *Monte Carlo Statistical Methods, 2nd ed.* Springer-Verlag, New York, 2004.
- [77] Jun S.-H. and A. Bouchard-Côté. Memory (and time) efficient sequential Monte Carlo. In T. Jebara and E.P. Xing, editors, *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, pages 514–522, 2014. URL <http://jmlr.org/proceedings/papers/v32/jun14.pdf>.
- [78] Jun S.-H., L. Wang, and A. Bouchard-Côté. Entangled Monte Carlo. In F. Pereira, C.J.C. Burges, L. Bottou, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems 25*, pages 2726–2734. 2012. URL <http://papers.nips.cc/paper/4648-entangled-monte-carlo.pdf>.
- [79] N. Schweizer. Non-asymptotic error bounds for sequential MCMC methods in multimodal settings. *arXiv preprint arXiv:1205.6733*, 2012.
- [80] N. Schweizer. Non-asymptotic error bounds for sequential MCMC and stability of Feynman–Kac propagators. *arXiv preprint arXiv:1204.2382*, 2012.
- [81] C. Sherlock, A.H. Thiery, G.O. Roberts, and J.S. Rosenthal. On the efficiency of pseudo-marginal random walk Metropolis algorithms. *The Annals of Statistics*, 43(1):238–275, 2015. URL <http://dx.doi.org/10.1214/14-AOS1278>.
- [82] V. Solo. Adaptive spectral factorization. *IEEE Transactions on Automatic Control*, 34(10):1047–1051, 1989.
- [83] G. Storvik. Particle filters for state-space models with the presence of unknown static parameters. *IEEE Transactions on Signal Processing*, 50:281–289, 2002.
- [84] Y.W. Teh, M.I. Jordan, M.J. Beal, and D.M. Blei. Hierarchical Dirichlet processes. *Journal of the American Statistical Association*, 101(476), 2006.

- [85] T. Toni, D. Welch, N. Strelkowa, A. Ipsen, and M.P.H. Stumpf. Approximate Bayesian computation scheme for parameter inference and model selection in dynamical systems. *Journal of the Royal Society Interface*, 6(31):187, 2009.
- [86] N. Tripuraneni, S. Gu, H. Ge, and Z. Ghahramani. A linear-time particle Gibbs sampler for infinite hidden Markov models. *arXiv preprint arXiv:1505.00428*, 2015.
- [87] C. Vergé, C. Dubarry, P. Del Moral, and E. Moulines. On parallel implementation of sequential Monte Carlo methods: the island particle model. *Statistics and Computing*, pages 1–18, 2013.
- [88] N. Whiteley. Sequential Monte Carlo samplers: error bounds and insensitivity to initial conditions. *Stochastic Analysis and Applications*, 30(5):774–798, 2012.
- [89] N. Whiteley. Stability properties of some particle filters. *The Annals of Applied Probability*, 23(6):2500–2537, 2013.
- [90] N. Whiteley, A. Lee, and K. Heine. On the role of interaction in sequential Monte Carlo algorithms. *Bernoulli*, to appear, 2015.
- [91] F. Wood, J.W. van de Meent, and V. Mansinghka. A new approach to probabilistic programming inference. In *Proceedings of the 17th International conference on Artificial Intelligence and Statistics*, pages 2–46, 2014.
- [92] C. Yau, O. Papaspiliopoulos, G. O. Roberts, and C. Holmes. Bayesian non-parametric hidden Markov models with applications in genomics. *Journal of the Royal Statistical Society: Series B*, 73(1):37–57, 2011. ISSN 1467-9868. URL <http://dx.doi.org/10.1111/j.1467-9868.2010.00756.x>.
- [93] Y. Zhou and A. Jasra. Biased online parameter inference for state-space models. *arXiv preprint arXiv:1503.00266*, 2015.
- [94] Y. Zhou, A.M. Johansen, and J.A.D. Aston. Towards automatic model comparison: an adaptive sequential Monte Carlo approach. *Journal of Computational and Graphical Statistics*, to appear, 2015.