Mathematical Biosciences 255 (2014) 21–32

Contents lists available at ScienceDirect

Mathematical Biosciences

journal homepage: www.elsevier.com/locate/mbs

Comparison of the performance of particle filter algorithms applied to tracking of a disease epidemic

Daniel M. Sheinson [a,*], Jarad Niemi [b], Wendy Meiring [a]

[a] Department of Statistics and Applied Probability, University of California, Santa Barbara, CA 93106, USA
[b] Department of Statistics, Iowa State University, Ames, IA 50011, USA

ARTICLE INFO

Article history:
Received 5 July 2013
Received in revised form 28 April 2014
Accepted 27 June 2014
Available online 9 July 2014

Keywords:
Bayesian estimation
Epidemiological modeling
State-space modeling
Sequential Monte Carlo
Particle filtering

ABSTRACT

We present general methodology for sequential inference in nonlinear stochastic state-space models to simultaneously estimate dynamic states and fixed parameters. We show that basic particle filters may fail due to degeneracy in fixed parameter estimation and suggest the use of a kernel density approximation to the filtered distribution of the fixed parameters to allow the fixed parameters to regenerate. In addition, we show that "seemingly" uninformative uniform priors on fixed parameters can affect posterior inferences and suggest the use of priors bounded only by the support of the parameter. We show the negative impact of using multinomial resampling and suggest the use of either stratified or residual resampling within the particle filter. As a motivating example, we use a model for tracking and prediction of a disease outbreak via a syndromic surveillance system. Finally, we use this improved particle filtering methodology to relax prior assumptions on model parameters yet still provide reasonable estimates for model parameters and disease states.

1. Introduction

State-space models are commonly used in the analysis of biological data [1–4]. In particular, these models are frequently used for disease outbreaks to simultaneously model the underlying disease dynamics and the observation process [5–7,3,8]. Together with syndromic surveillance systems [9–13], these models are used to identify emerging disease outbreaks [14], estimate their severity [6], and predict their duration [7]. Typically, the general form of these models may be reasonably assumed, but the models will have unknown fixed parameters that define the disease dynamics and the observation process for a particular outbreak. Based on similar previous outbreaks, the range and likely values for these fixed parameters may be available.

In statistical applications where prior knowledge or beliefs about unknown quantities are available, the Bayesian framework is often convenient for performing statistical analysis. Bayesian inference is conducted through the posterior distribution of any unknown quantities, obtained by updating prior information using observed data. However, the calculation of the posterior distribution in state-space models frequently involves complicated integrals without an explicit analytical form. The most common approach to approximating these posterior distributions is Markov chain Monte Carlo (MCMC) [15]. In a sequential context, e.g. syndromic surveillance, MCMC is inefficient due to the increase in computational cost incurred by the need for the entire MCMC to be rerun as each new observation arrives. Sequential Monte Carlo (SMC) – or particle filtering – methods enable on-line inference by updating the estimate of the posterior as new data become available. Furthermore, SMC methods can be flexible, general, easy to implement, and amenable to parallel computing. For a general introduction, please see [16,17].

Early SMC methods, including the bootstrap filter [18,19] and the auxiliary particle filter [20], assumed all fixed parameters were known. A key defining step in these filters is the use of resampling, which results in particles with low probability being eliminated and particles with high probability being duplicated. When all fixed parameters are known, these filters work remarkably well. In the presence of unknown fixed parameters, these filters suffer dramatically from a *degeneracy* issue due to the fixed parameters (being treated as dynamic states with degenerate evolutions) never being regenerated, and thus only a few distinct values for the fixed parameters remain after a few time points. To combat this degeneracy, a number of alternative approaches have been introduced. In

Abbreviations: BF, bootstrap filter; APF, auxiliary particle filter; KDPF, kernel density particle filter.

* Corresponding author. Tel.: +1 847 609 7824.
E-mail addresses: sheinson@pstat.ucsb.edu (D.M. Sheinson), niemi@iastate.edu (J. Niemi), meiring@pstat.ucsb.edu (W. Meiring).

http://dx.doi.org/10.1016/j.mbs.2014.06.018
0025-5564/

this paper, we focus on the *kernel density particle filter* [21] due to its wide applicability, ease of implementation, and good performance.

The rest of the article proceeds as follows. Section 2 contains a description of state-space models and sequential estimation. Section 3 describes a variety of particle filtering methodologies including the kernel density approach of [21]. In Section 4, we introduce a nonlinear dynamic model for a disease epidemic similar to the one used in Skvortsov and Ristic [3]. In Section 5, we apply the particle filtering methods described in Section 3 to the model described in Section 4. In Section 6, we benefit from the efficiency of more recent particle filtering methods to estimate a more complicated model. In Section 7, we conclude by reviewing more advanced particle filtering methods as well as the general performance and capabilities of SMC algorithms.

2. State-space models

State-space models are a general class of statistical models used for analysis of dynamic data and have been used extensively in modeling disease outbreaks [5,22,6,7,3]. State space models are constructed using an observation equation, $y_t \sim p_{y,t}(y_t|x_t,\theta)$, and a state evolution equation, $x_t \sim p_{x,t}(x_t|x_{t-1},\theta)$, where y_t is the observed response, x_t is a latent, dynamic state, and θ is an unknown fixed parameter, all of which could be vectors. The y_ts are assumed independent given x_t and θ, and the x_ts are assumed independent given x_{t-1} and θ. The distributions $p_{y,t}$ and $p_{x,t}$ are assumed known conditional on the values of θ and x_t in the observation equation and θ and x_{t-1} in the evolution equation. Depending on whether the observations and the states are continuous or discrete, the distributions themselves may be continuous or discrete. The distributions are typically assumed to only vary with x_t and θ, and therefore the t subscript is dropped. For simplicity, we also drop the x and y subscript and instead let the arguments make clear which distribution we are referring to. Thus, the general state-space model is

$$y_t \sim p(y_t|x_t,\theta) \quad x_t \sim p(x_t|x_{t-1},\theta).$$

A fully specified Bayesian model is obtained by also specifying the prior $p(x_0,\theta)$.

The dimension of x_t need not remain constant with respect to t. For instance, we could describe a process where x_t depends on the entire history of states up to t by letting $x_{t-1} = (x_1^*, x_2^*, \ldots, x_{t-1}^*)'$ and defining $x_t = (x_{t-1}, x_t^*)'$, where x_t^* is the new state generated at time t.

Special cases of these state-space models include hidden Markov models [23,24], where the state x_t has discrete support, and dynamic linear models (DLMs) [25], where each distribution is Gaussian whose mean is a linear function of the states and whose variance does not depend on the mean. The disease outbreak models discussed in Section 4 are specific cases of state-space models, but we introduce these models in generality here because the particle filtering methods discussed in Section 3 apply to any model with this form.

2.1. Sequential estimation

When data are collected sequentially, it is often of interest to determine the *filtered distribution*, the distribution of the current state and parameters conditional on the data observed up to that time. This distribution describes all of the available information up to time t about the current state of the system and any fixed parameters. It can be updated recursively using Bayes' rule:

$$p(x_t,\theta|y_{1:t}) \propto p(y_t|x_t,\theta)p(x_t,\theta|y_{1:t-1}) \tag{1}$$

where $y_{1:t} = (y_1, \ldots, y_t)$. Only in special cases can $p(x_t,\theta|y_{1:t})$ be evaluated analytically, e.g. in DLMs when θ is the observation variance [Section 4.3, 24]. When analytical tractability is not present, we turn to numerical methods including deterministic versions, e.g. the extended Kalman filter and the Gaussian sum filter [26], or Monte Carlo versions such as particle filters.

3. Particle filtering

Particle filtering is an SMC inferential technique based on repeated use of importance sampling. It aims to approximate the filtered distribution at time t through a weighted Monte Carlo realization from this distribution in terms of J particles, i.e.

$$p(x_t,\theta|y_{1:t}) \approx \sum_{j=1}^{J} w_t^{(j)} \delta_{\left(x_t^{(j)},\theta^{(j)}\right)} \tag{2}$$

where $\left(x_t^{(j)},\theta^{(j)}\right)$ is the location of the j^{th} particle at time t, $w_t^{(j)}$ is the weight of that particle with $\sum_{j=1}^{J} w_t^{(j)} = 1$, and δ is the Dirac delta function. A variety of SMC techniques have been developed to provide more efficient approximations to Eq. (1) in the sense that with the same computation time a better approximation is achieved. We now review three fundamental particle filtering techniques: the bootstrap filter, auxiliary particle filter, and kernel density particle filter. In Section 5, we compare the efficiency of these techniques in the syndromic surveillance context.

3.1. Bootstrap filter

The first successful version of particle filtering is known as the bootstrap filter (BF) [18,19]. Since this method and the auxiliary particle filter were developed for the situation when θ is known, we will (for the moment) drop θ from the notation. Given an approximation to the filtered distribution at time t as in Eq. (2), to obtain an approximation to the filtered distribution at time $t+1$, perform the following steps for each particle $j = 1, \ldots, J$:

1. Resample: sample an index $k \in \{1, \ldots, j, \ldots, J\}$ with associated probabilities $\left\{w_t^{(1)}, \ldots, w_t^{(j)}, \ldots, w_t^{(J)}\right\}$,
2. Propagate: sample $x_{t+1}^{(j)} \sim p\left(x_{t+1}\middle|x_t^{(k)}\right)$, and
3. Calculate weights and renormalize:

$$\tilde{w}_{t+1}^{(j)} = p\left(y_{t+1}\middle|x_{t+1}^{(j)}\right) \qquad w_{t+1}^{(j)} = \tilde{w}_{t+1}^{(j)} \Big/ \sum_{l=1}^{J} \tilde{w}_{t+1}^{(l)}.$$

This procedure can be applied recursively beginning with an initial set of weights $w_0^{(j)}$ and locations $x_0^{(j)}$ for all j, usually obtained by sampling from the prior with uniform weights.

3.2. Auxiliary particle filter

One problem that arises in implementing the BF is that $w_t^{(j)}$ will be small for particles where $p\left(y_t\middle|x_t^{(j)}\right)$ is small, and these particles will contribute little to the approximation to $p(x_t|y_{1:t})$. The auxiliary particle filter (APF) aims to mitigate this by anticipating which particles will have small weight using a look ahead strategy [20]. Given an approximation to the filtered distribution at time t as in Eq. (2), the APF approximates $p(x_{t+1}|y_{1:t+1})$ by the following:

1. For each particle j, calculate a point estimate of $x_{t+1}^{(j)}$ called $\mu_{t+1}^{(j)}$, e.g.

$$\mu_{t+1}^{(j)} = E\left(x_{t+1}\middle|x_t^{(j)}\right).$$