

Sequential Monte Carlo with Highly Informative Observations*

Pierre Del Moral[†] and Lawrence M. Murray[‡]

Abstract. We propose sequential Monte Carlo (SMC) methods for sampling the posterior distribution of state-space models under highly informative observation regimes, a situation in which standard SMC methods can perform poorly. A special case is simulating bridges between given initial and final values. The basic idea is to introduce a schedule of intermediate weighting and resampling times between observation times, which guide particles towards the final state. This can always be done for continuous-time models, and may be done for discrete-time models under sparse observation regimes; our main focus is on continuous-time diffusion processes. The methods are broadly applicable in that they support multivariate models with partial observation, do not require simulation of the backward transition (which is often unavailable), and, where possible, avoid pointwise evaluation of the forward transition. When simulating bridges, the last characteristic cannot be avoided entirely without concessions, and we suggest an ϵ -ball approach (reminiscent of approximate Bayesian computation) as a workaround. Compared to the bootstrap particle filter, the new methods deliver substantially reduced mean squared error in normalizing constant estimates, even after accounting for execution time. The methods are demonstrated for state estimation with two toy examples, and for parameter estimation (within a particle marginal Metropolis–Hastings sampler) with three applied examples in econometrics, epidemiology, and marine biogeochemistry.

Key words. particle filter, state-space model, Bayesian statistics, approximate Bayesian computation, continuous time, stochastic differential equation, Brownian bridge, diffusion bridge, pinned diffusion

AMS subject classifications. 65C05, 65C30, 65C35

DOI. 10.1137/15M1011214

1. Introduction. Consider the multivariate and continuous-time Markov process $X(t) \in \mathbb{R}^d$ and parameters Θ . For a sequence of times t_0, \dots, t_n we write $X_{0:n} \equiv \{X_0, \dots, X_n\} \equiv \{X(t_0), \dots, X(t_n)\}$, and adopt the convention that uppercase symbols denote random variables, with matching lowercase symbols realizations of them. We have $X_k \sim p(dx_k | x_{k-1}, \theta)$. The process may be observed indirectly via some $Y_n \sim p(dy_n | x_n, \theta)$, or directly as some given initial value x_0 and final value x_n . This setup admits continuous-time models, where n can be made arbitrarily large and thus t_1, \dots, t_{n-1} arbitrarily dense, and discrete-time models with sparse observation, where n cannot be made arbitrarily large, but where times t_1, \dots, t_{n-1} are unobserved.

In the case of indirect observation, the problem of interest is to simulate $X_{0:n} \sim p(dx_{0:n} | y_n, \theta)$ and perhaps estimate the normalizing constant (marginal likelihood) $p(y_n | \theta)$. We are particularly interested in the case when the observation is in some sense highly informative on $X_{0:n}$. This might arise when highly accurate measurements are taken in controlled settings,

*Received by the editors March 5, 2015; accepted for publication (in revised form) July 16, 2015; published electronically October 1, 2015.

<http://www.siam.org/journals/juq/3/M101121.html>

[†]University of New South Wales, Sydney, Australia (p.del-moral@unsw.edu.au).

[‡]Department of Statistics, University of Oxford, Oxford, OX1 3TG, UK (murray@stats.ox.ac.uk).

or where observations are relatively sparse in time. It will be adequate in this work for the meaning of “highly informative” to remain qualitative—vague even—but if it were to be quantified it might be defined as some large divergence (e.g., Kullback–Leibler) of $p(dx_n | \theta)$ from $p(dx_n | y_n, \theta)$.

In the case of direct observation, the problem of interest is to simulate *bridges* $X_{1:n-1} \sim p(dx_{1:n-1} | x_0, x_n, \theta)$ and perhaps estimate the normalizing constant (transition density) $p(x_n | x_0, \theta)$. This might be seen as the special case of indirect observation with $p(y_n | x_n) = \delta(y_n - x_n)$, where δ is the Dirac δ function. We motivate the approach with this special case, and return to the more general case later.

The simulation of bridges is generally regarded as a difficult problem, and there is a wealth of literature in the area. Most approaches, while not necessarily limited to this area, begin with a diffusion process satisfying the Itô stochastic differential equation (SDE)

$$dX(t) = a(X(t), t, \theta) dt + B(X(t), t, \theta) dW(t),$$

where $a(X(t), t, \theta)$ is the drift vector, $B(X(t), t, \theta)$ is the diffusion matrix, and $W(t)$ is a vector of standard Wiener processes. One group of methods proceeds with a schedule of times t_0, \dots, t_n , equispaced at a step size sufficiently small that a locally linear–Gaussian Euler–Maruyama [25, sect. 9.1] discretization of the original dynamics makes a credible approximation:

$$p(x_k | x_{k-1}, \theta) := \phi(x_{k-1} + a(x_{k-1}, t_{k-1}, \theta)\Delta t_k, B(x_{k-1}, t_{k-1}, \theta)B(x_{k-1}, t_{k-1}, \theta)^\top \Delta t_k).$$

Here, $k = 1, \dots, n$, $\Delta t_k = t_k - t_{k-1}$, and $\phi(\mu, \Sigma)$ is the probability density function of the normal distribution with mean vector μ and covariance matrix Σ . Because of the convenience of this closed form for $p(x_k | x_{k-1}, \theta)$, the discretized version may be adopted in place of the continuous version (as an a priori approximation) before proceeding with inference. This facilitates various importance sampling (e.g., [36, 16]), sequential Monte Carlo (SMC, e.g., [27]), and Markov chain Monte Carlo (MCMC, e.g., [40, 17, 18, 21]) methods for simulating bridges.

Caution is warranted around the use of Euler–Maruyama, however. While the discretization provides a convenient approximation for pointwise evaluation of $p(x_k | x_{k-1}, \theta)$, it can be unstable for simulation unless the step size is very small—too small for efficient computation. In such cases higher-order schemes, such as the Milstein and semi-implicit schemes [25], or those of the Runge–Kutta family, are to be preferred, and indeed may be essential. For any given step size Δt , these higher-order schemes have stability regions at least as large as that of Euler–Maruyama. For particularly stiff problems, implicit schemes may also need to be considered.

Rather than discretizing to yield a closed-form transition density, a closed-form Radon–Nikodym derivative can be derived in certain conditions. This leads to another family of methods for simulating bridges (e.g., [10, 14, 4, 42]). The exact algorithm [7] is yet another alternative, and has the advantage of not introducing discretization error, but its requirements restrict the set of models to which it can be applied, especially the set of multivariate models.

The estimation of the normalizing constant $p(x_n | x_0, \theta)$ usually drops out as a straightforward expectation in importance sampling methods, including SMC. It is more difficult

with MCMC. A number of works have focused more closely on this, usually in the context of obtaining normalizing constant estimates for parameter estimation (e.g., [19, 20, 44]).

The methods proposed in this work are most similar to the SMC methods of Del Moral and Garnier [12] and Lin, Chen, and Mykland [27]. The former considers the simulation of rare events, not bridges, albeit with similar mechanisms. The latter does consider the simulation of bridges, but the implementation has some limitations, which this work seeks to ameliorate. In particular, Lin, Chen, and Mykland [27] require that pilot samples can be initialized at time t_n and simulated backwards in time to guide those samples being simulated forwards in time. This is problematic in cases when the backwards transition cannot be discretized in a numerically stable way. In contrast, the proposed methods do not require simulation of the backwards transition. In addition, Lin, Chen, and Mykland [27] do not support indirect observation, and in the case of multivariate models and direct observation, does not support partial observation of the state vector x_n . The proposed methods accommodate both of these cases. In the context of indirect observation, the methods are similar to the auxiliary particle filter [38], but perform lookahead steps at multiple intermediate times between observations, rather than at observation times only.

The remainder of this work describes the proposed methods for sampling from state-space models with highly informative observations, including the special case of simulating bridges. The methods also permit estimation of normalizing constants. The basic idea is to simulate particles forward in time using only the prior $p(dx_k | x_{k-1}, \theta)$, discretized with a higher-order scheme, but also to introduce additional weighting and resampling steps at intermediate times to guide particles towards the final state. The methods are simple to apply, with the following properties that make them useful for a broad range of problems:

1. They work in a multivariate setting.
2. They support partial observation of the state vector x_n .
3. They do not require that the backward transition can be simulated.
4. They support higher-order discretizations of the forward state process than that of Euler–Maruyama.
5. They require only that the forward transition can be simulated and not that its probability density function $p(x_k | x_{k-1}, \theta)$ can be evaluated pointwise. This comes with the caveat that in the special case of simulating bridges, a workaround is needed to approximately evaluate, or avoid the evaluation of, $p(x_n | x_{n-1}, \theta)$. In one example we use an Euler–Maruyama discretization to approximately evaluate, but not simulate, $p(x_n | x_{n-1}, \theta)$. In another we concede an ϵ -ball around x_n , where ϵ is commensurate with the discretization error already surrendered by the numerical integrator. The latter strategy resembles a simple approximate Bayesian computation (ABC) algorithm [5].

Three methods are introduced in this work, all similar but tailored for slightly different circumstances. Section 2 provides the formal background to establish the methods. Section 3 provides pseudocode for their implementation and discusses design of the required weight functions. Section 4 provides empirical results for the methods on two toy examples and three applications in econometrics, epidemiology, and marine biogeochemistry. Section 5 draws these results together and reports on experiences in tuning the methods. Section 6 concludes.

2. Methods. We derive three methods in this section, all quite related, but for different circumstances of the model and data. References to the parameters, Θ , are omitted throughout for brevity. Likewise, observation of the final value x_n may be full or partial; the methods support both cases, but for simplicity no notational distinction is made.

First note that

$$\begin{aligned} p(x_{1:n-1} | x_0, x_n) &= \frac{p(x_n | x_{n-1})p(x_{1:n-1} | x_0)}{p(x_n | x_0)} \\ &\propto p(x_n | x_{n-1})p(x_{1:n-1} | x_0). \end{aligned}$$

This forms the basis of the importance sampling method of Pedersen [36], proposing from $p(dx_{1:n-1} | x_0)$ and weighting with $p(x_n | x_{n-1})$. The use of the prior as the proposal in this way is myopic of x_n , and while often workable, the approach can lead to excessive variance in importance weights and subsequent computational expense. In Durham and Gallant [16], an alternative proposal is suggested, adjusting the drift term of the SDE with a linear component to draw the process towards x_n . The transition densities $p(x_k | x_{k-1})$ must then be estimated in order to compute weights, and the low-order Euler–Maruyama discretization is employed to achieve this. One, perhaps overlooked, advantage of Pedersen [36] is that $p(x_k | x_{k-1})$ does not need to be evaluated pointwise except for the weight at $k = n$. This means that for the purposes of simulation, higher-order discretization schemes can be admitted. The methods of this work benefit from the same property.

Further observe that

$$(2.1) \quad p(x_{1:n-1} | x_0) = \prod_{k=1}^{n-1} p(x_k | x_{k-1})$$

and

$$(2.2) \quad \frac{p(x_n | x_{n-1})}{p(x_n | x_0)} = \prod_{k=1}^{n-1} \frac{p(x_n | x_k)}{p(x_n | x_{k-1})},$$

and define

$$G_k(x_{k-1:k}) := \frac{p(x_n | x_k)}{p(x_n | x_{k-1})}.$$

One can then incrementally simulate $p(dx_k | x_{k-1})$ and weight with $G_k(x_{k-1:k})$, for $k = 1, \dots, n-1$. The basis of an SMC method is precisely this, maintaining a population of samples (particles) and introducing a selection mechanism to resample from amongst them, according to their weights, at each increment. Lin, Chen, and Mykland [27] does this; the development below follows a similar path, but we will ultimately suggest a different implementation, and extend the idea from the special case of sampling bridges to the more general case of sampling under highly informative observations.

Proposition 2.1. *For any bounded function f on $(\mathbb{R}^d)^{n-1}$, we have*

$$\mathbb{E}[f(X_{1:n-1}) | x_0, x_n] = \mathbb{E} \left[f(X_{1:n-1}) \prod_{k=1}^{n-1} G_k(X_{k-1:k}) \middle| x_0 \right]$$

with the weight functions

$$G_k(x_{k-1:k}) := \frac{p(x_n | x_k)}{p(x_n | x_{k-1})}.$$

Proof. The conditional density of the random path $X_{1:n-1}$ given initial state $X_0 = x_0$ and final value $X_n = x_n$ is

$$\begin{aligned} p(x_{1:n-1} | x_0, x_n) &= \frac{p(x_n | x_{n-1})p(x_{1:n-1} | x_0)}{p(x_n | x_0)} \\ &= \left\{ \prod_{k=1}^{n-1} \frac{p(x_n | x_k)}{p(x_n | x_{k-1})} \right\} p(x_{1:n-1} | x_0) \\ &= \left\{ \prod_{k=1}^{n-1} G_k(x_{k-1:k}) \right\} p(x_{1:n-1} | x_0). \end{aligned}$$

Thus for any bounded function f on $(\mathbb{R}^d)^{n-1}$ we have

$$\begin{aligned} \mathbb{E}[f(X_{1:n-1}) | x_0, x_n] &= \int_{\mathbb{R}^{d(n-1)}} f(x_{1:n-1}) p(x_{1:n-1} | x_0, x_n) dx_{1:n-1} \\ &= \int_{\mathbb{R}^{d(n-1)}} f(x_{1:n-1}) \left\{ \prod_{k=1}^{n-1} G_k(x_{k-1:k}) \right\} p(x_{1:n-1} | x_0) dx_{1:n-1} \\ &= \mathbb{E} \left[f(X_{1:n-1}) \prod_{k=1}^{n-1} G_k(X_{k-1:k}) \middle| x_0 \right]. \quad \blacksquare \end{aligned}$$

Of course, rarely will the weights $G_k(x_{k-1:k})$ be computable in practice; Proposition 2.1 is conceptually appealing, however, and we can try to imitate it in other circumstances. To do this, we introduce arbitrary weighting functions that facilitate an SMC algorithm.

Proposition 2.2. For any bounded function f on $(\mathbb{R}^d)^{n-1}$, we have

$$\mathbb{E}[f(X_{1:n-1}) | x_0, x_n] = \frac{\mathbb{E} \left[f(X_{1:n-1}) p(x_n | X_{n-1}) \frac{q(x_n | x_0)}{q(x_n | X_{n-1})} \prod_{k=1}^{n-1} H_k(X_{k-1:k}) \middle| x_0 \right]}{\mathbb{E} \left[p(x_n | X_{n-1}) \frac{q(x_n | x_0)}{q(x_n | X_{n-1})} \prod_{k=1}^{n-1} H_k(X_{k-1:k}) \middle| x_0 \right]}$$

with the weight functions

$$H_k(x_{k-1:k}) := \frac{q(x_n | x_k)}{q(x_n | x_{k-1})}$$

and chosen positive functions $q(x_n | x_k)$ for $k = 0, \dots, N-1$.

Proof. Observe that

$$p(x_n | x_0) = \mathbb{E}[p(x_n | X_{n-1}) | x_0]$$

so that

$$\begin{aligned} \mathbb{E}[f(X_{1:n-1}) | x_0, x_n] &= \mathbb{E} \left[f(X_{1:n-1}) \frac{p(x_n | X_{n-1})}{p(x_n | x_0)} \middle| x_0 \right] \\ &= \frac{\mathbb{E}[f(X_{1:n-1}) p(x_n | X_{n-1}) | x_0]}{p(x_n | x_0)} \\ &= \frac{\mathbb{E}[f(X_{1:n-1}) p(x_n | X_{n-1}) | x_0]}{\mathbb{E}[p(x_n | X_{n-1}) | x_0]}. \end{aligned}$$

We introduce

$$\frac{q(x_n | x_0)}{q(x_n | x_{n-1})} \prod_{k=1}^{n-1} H_k(x_{k-1:k}) = 1.$$

We then have

$$\mathbb{E}[f(X_{1:n-1}) | x_0, x_n] = \frac{\mathbb{E}\left[f(X_{1:n-1})p(x_n | X_{n-1})\frac{q(x_n | x_0)}{q(x_n | X_{n-1})} \prod_{k=1}^{n-1} H_k(X_{k-1:k}) \middle| x_0\right]}{\mathbb{E}\left[p(x_n | X_{n-1})\frac{q(x_n | x_0)}{q(x_n | X_{n-1})} \prod_{k=1}^{n-1} H_k(X_{k-1:k}) \middle| x_0\right]}.$$

The results are easily adapted to the case where the state is not observed exactly because of some noise. We introduce the random variable $Y_n \sim p(dy_n | x_n)$, which will be observed in place of X_n , and a prior distribution over the starting state, $X_0 \sim p(dx_0)$.

Proposition 2.3. *For any bounded function f on $\mathbb{R}^{d(n+1)}$, we have*

$$\mathbb{E}[f(X_{0:n}) | y_n] = \frac{\mathbb{E}\left[f(X_{0:n})p(X_0)p(y_n | X_n)\frac{r(y_n | X_0)}{r(y_n | X_n)} \prod_{k=1}^n J_k(X_{k-1:k})\right]}{\mathbb{E}\left[p(X_0)p(y_n | X_n)\frac{r(y_n | X_0)}{r(y_n | X_n)} \prod_{k=1}^n J_k(X_{k-1:k})\right]}$$

with the weight functions

$$J_k(x_{k-1:k}) := \frac{r(y_n | x_k)}{r(y_n | x_{k-1})}$$

and chosen positive functions $r(y_n | x_k)$ for $k = 0, \dots, N$.

Proof. Observe that

$$p(y_n) = \mathbb{E}[p(y_n | X_n)]$$

so that

$$\begin{aligned} \mathbb{E}[f(X_{0:n}) | y_n] &= \mathbb{E}\left[f(X_{0:n})\frac{p(y_n | X_n)}{p(y_n)}\right] \\ &= \frac{\mathbb{E}[f(X_{0:n})p(y_n | X_n)]}{p(y_n)} \\ &= \frac{\mathbb{E}[f(X_{0:n})p(y_n | X_n)]}{\mathbb{E}[p(y_n | X_n)]}. \end{aligned}$$

We introduce

$$\frac{r(y_n | x_0)}{r(y_n | x_n)} \prod_{k=1}^n J_k(x_{k-1:k}) = 1.$$

We then have

$$\mathbb{E}[f(X_{0:n}) | y_n] = \frac{\mathbb{E}\left[f(X_{0:n})p(X_0)p(y_n | X_n)\frac{r(y_n | X_0)}{r(y_n | X_n)} \prod_{k=1}^n J_k(X_{k-1:k})\right]}{\mathbb{E}\left[p(X_0)p(y_n | X_n)\frac{r(y_n | X_0)}{r(y_n | X_n)} \prod_{k=1}^n J_k(X_{k-1:k})\right]}.$$

We expect this modification to lead to an SMC algorithm that is particularly useful when observations are highly informative.

3. Implementation. The recursive structure of the weight functions in Propositions 2.1–2.3 and the Markov property of the process $X(t)$ facilitate SMC algorithms to compute the expectations of interest. Such algorithms propagate, weigh, and resample a population of N particles. We present Algorithms 3.1–3.3 as pseudocode, corresponding to Propositions 2.1–2.3, respectively. Where a superscript i appears on the left-hand side of an assignment in these algorithms, the intended interpretation is “for all $i \in \{1, \dots, N\}$.” The left arrow notation (\leftarrow) denotes assignment of the value on the right to the variable on the left, while the tilde notation (\sim) denotes assignment of a draw from the distribution on the right to the variable on the left. We start with Algorithm 3.1, corresponding to Proposition 2.1.

Algorithm 3.1.

```

 $x_0^i \leftarrow x_0$  // initialize
 $w_0^i \leftarrow 1/N$ 
for  $k = 1, \dots, n - 1$ 
  if resampling is triggered
     $a_k^i \sim R(da_k | w_{k-1}^{1:N})$  // select
     $w_k^i \leftarrow 1/N$ 
  else
     $a_k^i \leftarrow i$ 
     $w_k^i \leftarrow w_{k-1}^i / \sum_j w_{k-1}^j$ 
     $x_k^i \sim p(dx_k | x_{k-1}^{a_k^i})$  // propagate
     $w_k^i \leftarrow p(x_n | x_k^i) \cdot w_k^i / w_{k-1}^{a_k^i}$  // weight.

```

Algorithm 3.2, corresponding to Proposition 2.2, is given below. It assumes that $q(x_n | x_{n-1}) := p(x_n | x_{n-1})$. Note that Algorithm 3.1 is just the special case of Algorithm 3.2 where $q(x_n | x_k) := p(x_n | x_k)$.

Algorithm 3.2.

```

 $x_0^i \leftarrow x_0$  // initialize
 $w_0^i \leftarrow 1/N$ 
for  $k = 1, \dots, n - 1$ 
  if resampling is triggered
     $a_k^i \sim R(da_k | w_{k-1}^{1:N})$  // select
     $w_k^i \leftarrow 1/N$ 
  else
     $a_k^i \leftarrow i$ 
     $w_k^i \leftarrow w_{k-1}^i / \sum_j w_{k-1}^j$ 
     $x_k^i \sim p(dx_k | x_{k-1}^{a_k^i})$  // propagate
     $w_k^i \leftarrow q(x_n | x_k^i) \cdot w_k^i / w_{k-1}^{a_k^i}$  // weight.

```

At the conclusion of Algorithm 3.1 or 3.2, let $b_{n-1}^i = i$ and, recursively, $b_k^i = a_{k+1}^{b_{k+1}^i}$. The indices b_k^i then establish ancestral lines:

$$x_0 \rightarrow x_1^{b_1^i} \rightarrow x_2^{b_2^i} \rightarrow \dots \rightarrow x_{n-2}^{b_{n-2}^i} \rightarrow x_{n-1}^{b_{n-1}^i} \rightarrow x_n.$$

Because SMC methods are a particular case of mean field particle methods [11], these ancestral

lines may be used to compute expectations of the forms that appear in Propositions 2.1 and 2.2:

$$\frac{\sum_{i=1}^N w_{n-1}^i f(x_1^{b_i^1}, \dots, x_{n-1}^{b_i^{n-1}})}{\sum_{i=1}^N w_{n-1}^i} \xrightarrow{N \uparrow \infty} \mathbb{E}[f(X_{1:n-1}) | x_0, x_n].$$

The denominator on the left is also an estimate of the normalizing constant:

$$(3.1) \quad \sum_{i=1}^N w_{n-1}^i \xrightarrow{N \uparrow \infty} \mathbb{E}[p(x_n | X_{n-1}) | x_0] \\ = p(x_n | x_0).$$

Note that as Algorithms 3.1 and 3.2 normalize the weights $w_{n-1}^{1:N}$ after the selection step but before the weighting step, no factor of $1/N$ appears outside the summation in (3.1).

It is often the case that the transition density $p(x_n | x_{n-1})$ does not have a convenient closed form for pointwise evaluation, so that the last line of Algorithm 3.2 cannot be evaluated. In such cases one of two approaches might be considered. In the first approach, the sequence of times t_1, \dots, t_{n-1} might be set so that the last interval, $\Delta t_n = t_n - t_{n-1}$, is sufficiently small for the Euler–Maruyama approximation of $p(x_n | x_{n-1})$ to be credible. Because this last transition is evaluated but not simulated (much less simulated repeatedly with accumulating error), the stability issues of the Euler–Maruyama discretization will not manifest. In the second approach, an observation model might be constructed with an ϵ -ball around x_n , where ϵ is commensurate with the discretization error already inherent in the numerical integrator. We do this in the SIR example of section 4.

This second approach yields an algorithm resembling ABC [5]. In a simple ABC algorithm, one would simulate a path $x'_{1:n} \sim p(dx_{1:n} | x_0)$ and accept it if $\rho(x'_n, x_n) \leq \epsilon$ for some distance function ρ and error threshold ϵ . The SMC component of the proposed method marginalizes over multiple such paths. If we define the unnormalized density

$$p(x_n | x_{n-1}) := \begin{cases} 1 & \text{if } \rho(x'_n, x_n) \leq \epsilon, \\ 0 & \text{otherwise,} \end{cases}$$

then the estimate of the normalizing constant (3.1) is also an estimate of the acceptance probability of θ' . It is worth stressing that the SMC component in this case is used in a very different way than ABC SMC methods in the spirit of, e.g., Sisson, Fan, and Tanaka [43], Beaumont et al. [6], Del Moral, Doucet, and Jasra [13], and Peters, Fan, and Sisson [37]. In these works, SMC is used over parameters; here it is used over the state. It could be coupled with MCMC (as in particle MCMC, Andrieu, Doucet, and Holenstein [1]) or another level of SMC (as in SMC², Chopin, Jacob, and Papaspiliopoulos [9]) for parameter estimation, however.

Algorithm 3.3, corresponding to Proposition 2.3, is a slight variation on Algorithm 3.2, as x_0 and x_n are no longer fixed and y_n is introduced. It assumes, sensibly, that $r(y_n | x_n) := p(y_n | x_n)$.

Algorithm 3.3.

```

 $x_0^i \sim p(dx_0)$  // initialize
 $w_0^i \leftarrow 1/N$ 
for  $k = 1, \dots, n$ 
    if resampling is triggered
         $a_k^i \sim R(da_k | w_{k-1}^{1:N})$  // select
         $w_k^i \leftarrow 1/N$ 
    else
         $a_k^i \leftarrow i$ 
         $w_k^i \leftarrow w_{k-1}^i / \sum_j w_{k-1}^j$ 
     $x_k^i \sim p(dx_k | x_{k-1}^{a_k^i})$  // propagate
     $w_k^i \leftarrow r(y_n | x_k^i) \cdot w_k^i / w_{k-1}^{a_k^i}$  // weight.

```

At the conclusion of Algorithm 3.3, the indices b_k^i , defined as before, establish ancestral lines:

$$x_0^{b_0^i} \rightarrow x_1^{b_1^i} \rightarrow x_2^{b_2^i} \rightarrow \dots \rightarrow x_{n-2}^{b_{n-2}^i} \rightarrow x_{n-1}^{b_{n-1}^i} \rightarrow x_n^{b_n^i}.$$

These may be used to compute expectations of the form that appears in Proposition 2.3:

$$\frac{\sum_{i=1}^N w_n^i f(x_0^{b_0^i}, \dots, x_n^{b_n^i})}{\sum_{i=1}^N w_n^i} \xrightarrow{N \uparrow \infty} \mathbb{E}[f(X_{0:n}) | y_n].$$

The denominator on the left is also an estimate of the normalizing constant:

$$(3.2) \quad \sum_{i=1}^N w_n^i \xrightarrow{N \uparrow \infty} \mathbb{E}[p(y_n | X_n)] = p(y_n).$$

Note that as Algorithm 3.3 normalizes the weights $w_n^{1:N}$ after the selection step but before the weighting step, no factor of $1/N$ appears outside the summation in (3.2).

Algorithm 3.3 treats the case where there is only a single observation, at time t_n . This is straightforwardly extended to a time series of observations, where the algorithm is repeated, but the first two lines removed from the second and subsequent iterations; the current particles and their weights are maintained instead. This is then a particle filter, with the addition of intermediate times between observations where additional weighting and resampling is performed to guide particles towards the next state.

3.1. Intermediate weighting. What remains is the selection of appropriate functions q and r in Algorithms 3.2 and 3.3. For good performance, we should seek $q(x_n | x_k) \approx p(x_n | x_k)$, so that Algorithm 3.2 approximates Algorithm 3.1 as closely as possible, and $r(y_n | x_k) \approx p(y_n | x_k)$, much like the use of lookahead strategies [28] for stage one weights in an auxiliary particle filter [38]. As $q(x_n | x_k)$ and $r(y_n | x_k)$ play a role similar to the proposal distribution in importance sampling, we should also prefer that their tails are not too tight with respect to $p(x_n | x_k)$ or $p(y_n | x_k)$, respectively. For reasons of computational expediency, we suppose that the weight functions are to be selected a priori. Choosing some parametric form, we

might choose either to fit the function to simulations of the prior model, or to the data set. We utilize both approaches in the examples of section 4.

The implementation in Lin, Chen, and Mykland [27] uses a kernel density estimate of $p(x_k | x_n) \propto p(x_n | x_k)p(x_k)$, obtained by propagating pilot particles backwards from time t_n to time t_k , initializing each at x_n . This is problematic for the constraints we have set for ourselves: it requires that x_n is fully observed in order to initialize each particle, it does not support indirect observation, and we do not wish to assume that the backwards transition can be simulated in a numerically stable way.

For diffusion processes, discretizations that provide a closed-form transition density may be useful. As even very approximate weight functions may have some utility, an Euler–Maruyama discretization might prove useful. More sophisticated Gaussian approximations such as a linear noise approximation [45, p. 258] may also be useful, if more computationally expensive.

We have found that a generally useful approach is to fit a Gaussian process to each observed time series and then construct weight functions based on these:

$$X(t) \sim \mathcal{GP}(\mu(t), C(\Delta t)),$$

with mean function $\mu(t)$ and covariance function $C(\Delta t)$. This affords a great deal of flexibility in the design of weight functions, accommodating arbitrary mean functions to capture non-linear drifts, and a variety of covariance functions to capture local behavior and smoothness. In the examples of section 4 it has not been necessary to be too clever to obtain good results; the mean function is always $\mu(t) = 0$ and the covariance function of a squared exponential form, parameterized by α and β :

$$C(\Delta t) = \alpha \exp\left(-\frac{1}{2\beta}(\Delta t)^2\right).$$

The parameters α and β are set to their maximum likelihood estimates, obtained offline. One can, of course, imagine more sophisticated mean and covariance functions—Gaussian processes being very flexible in this regard—but we have found this simple choice adequate for the examples here. The functions $q(x_n | x_k)$ are then constructed by conditioning the Gaussian process on the current state, taking

$$q(x_n | x_k) := \phi(\mu_k, \sigma_k^2)$$

with

$$\begin{aligned}\mu_k &= \frac{C(t_n - t_k)}{\alpha} x_k, \\ \sigma_k^2 &= \alpha - \frac{C(t_n - t_k)^2}{\alpha}.\end{aligned}$$

Conditioning on the current state only, and not the full state history, is a computational concession, preserving linear complexity in the number of particles N . We have found this sufficient for the examples in this work. If necessary, one might consider conditioning on some

fixed number of previous states, preserving the same linear complexity. In the case of indirect observation, a (possibly approximate) Gaussian observation model of

$$Y(t) \sim \mathcal{N}(X(t), \varsigma^2(t)),$$

for some variance $\varsigma^2(t)$, would suggest weight functions of

$$r(y_n | x_k) := \phi(\mu_k, \sigma_k^2 + \varsigma_n^2).$$

We may be concerned that these tight-tailed Gaussians are too narrow as weight functions, and may be too aggressive in particle selection. A simple precaution is to inflate the variance by some constant factor, and we do this in the examples of section 4. One might also consider heavier-tailed functions such as the Student t . Another limitation of the Gaussian process formulation is that it may be inadequate for capturing multimodal transition densities. This occurs in the periodic drift example of section 4, and we propose a bespoke function in that case.

3.2. Intermediate resampling. The intermediate resampling steps deserves some consideration. In the first instance, they introduce additional computational expense. In the second they may introduce additional variance in normalizing constant estimates [39, 26]. In the case of computational cost, the numerical integration of diffusion processes (required to propagate particles forward) will typically be much more expensive than the steps required to resample. We might assume, then, that the additional resampling adds little to the overall cost. At any rate an adaptive resampling scheme mitigates both issues.

A simple adaptive resampling scheme is based on the *effective sample size* (ESS), which for the weight vector $w_k^{1:N}$ at time t_k is [29]

$$\text{ESS}(w_k^{1:N}) = \frac{\left(\sum_i w_k^i\right)^2}{\sum_i (w_k^i)^2}.$$

Resampling is then only triggered if this quantity falls below some threshold. We do this in the experimental results of section 4, and find that the net effect of the additional resampling steps is beneficial.

When an adaptive scheme such as this is used, the increase in variance should be constant with respect to the number of intermediate times. Intuitively, this is clear from (2.2): the accumulated weight at some time t_k is always $p(x_n | x_k)/p(x_n | x_0)$, regardless of the preceding time schedule t_1, \dots, t_{k-1} . Rather than determining the accumulated weight, the time schedule determines the times at which resampling should be considered. Note that if resampling is never triggered, all the additional weights cancel. Algorithm 3.2 then becomes the method of Pedersen [36], while Algorithm 3.3, iterated, becomes the bootstrap particle filter.

4. Experiments. We use five different examples to demonstrate the methods.

- OU:** a linear–Gaussian Ornstein–Uhlenbeck process fit to simulated data, without parameter estimation (cf. [44]),
- FFR:** a linear–Gaussian Ornstein–Uhlenbeck process fit to Federal Funds Rate data, with parameter estimation (cf. [3]),

PD: a nonlinear periodic drift process fit to simulated data, without parameter estimation (cf. [7, 27]),

SIR: a multivariate and nonlinear susceptible/infected/recovered compartmental model used in epidemiology, fit to influenza data from a boarding school [2], with parameter estimation (cf. [41]), and

NPZD: a multivariate and nonlinear nutrient/phytoplankton/zooplankton/detritus model [35] used in marine biogeochemistry, fit to Ocean Station P data [31], with parameter estimation (cf. [35, 34]).

The OU and PD examples are toy studies used to illustrate the methods, while the FFR, SIR, and NPZD examples are applied problems using real data sets. The SIR example has additional interest for the ABC-like approach taken.

Experiments are conducted using the LibBi software [33], in which the methods have been implemented under the name *bridge particle filter*. We use this name henceforth. Each example is available as a separate LibBi package, available from www.libbi.org. The bootstrap particle filter, as implemented in LibBi, is used for comparison, noting that with no resampling at intermediate times, it reduces to the method of Pedersen [36].

Configurations for all experiments are summarized in Table 1 and detailed in the text. Using LibBi, it is straightforward to run both the bootstrap and bridge particle filters across multiple threads on a central processing unit (CPU), with or without the use of SSE vector instructions, or on a graphics processing unit (GPU). Table 1 also documents the chosen hardware configuration for each example, chosen for fastest execution time after some pilot runs.

To compare methods, we use a number of metrics based on normalizing constant estimates, which are further scaled by execution time for a fair computational comparison. For a set of Z (see Table 1 for specifics) normalizing constant estimates, $z^{1:Z}$, obtained after corresponding execution times $t^{1:Z}$, we define the following metrics:

$$(4.1) \quad \text{MSE}(\log z^{1:Z})^{-1} \cdot \text{Mean}(t^{1:Z})^{-1},$$

$$(4.2) \quad \text{ESS}(z^{1:Z}) \cdot \text{Mean}(t^{1:Z})^{-1},$$

$$(4.3) \quad \text{CAR}(z^{1:Z}) \cdot \text{Mean}(t^{1:Z})^{-1}.$$

Here, $\text{Mean}(t^{1:Z})$ is the sample mean of $t^{1:Z}$. $\text{MSE}(z^{1:Z})$ is the mean-squared error of $z^{1:Z}$:

$$\text{MSE}(z^{1:Z}) = \frac{1}{Z} \sum_{i=1}^Z (z^i - z^*)^2,$$

where z^* is the true normalizing constant. If z^* is not known, the best available estimate is substituted as its “true” value. This will be the estimate obtained from a bootstrap particle filter¹ using a great many particles (see Table 1 for specifics). $\text{ESS}(z^{1:Z})$ is the effective sample size of $z^{1:Z}$ [29]:

$$\text{ESS}(z^{1:Z}) = \frac{\left(\sum_{i=1}^Z z^i\right)^2}{\sum_{i=1}^Z (z^i)^2}.$$

¹Our results indicate that the bridge particle filter should typically give a better estimate, but we avoid basing the truth on the method to be validated.

Table 1
Experimental configurations and some results for all examples.

Model	OU	FFR	PD	SIR	NPZD
Simulation time step	0.01	0.01	0.075	Adaptive	Adaptive
Bridge time step	0.1	0.1	1	0.01	1
Bridge type	Exact	Exact	Parametric	\mathcal{GP}	\mathcal{GP}
Data set					
N ^o observations	100	300	100	14	227
Observation time step	1	1	30	1	Irregular
Normalizing constant experiments					
N ^o data sets	16	1	16	-	1
N ^o parameter sets	1	16	1	-	16
N ^o particles (N)	$2^5, \dots, 2^{10}$	$2^5, \dots, 2^{10}$	$2^5, \dots, 2^{10}$	-	$2^5, \dots, 2^{10}$
Total N ^o experiments	96	96	96	-	96
N ^o repetitions on each experiment (Z)	2^{12}	2^{12}	2^{12}	-	2^{10}
N ^o particles (N) for “true” log-likelihood	-	-	2^{20}	-	2^{20}
Parameter estimation experiments					
N ^o MCMC steps after burn-in (N_θ)	-	1×10^5	-	5×10^4	1×10^5
N ^o particles (N)	-	2^8	-	2^{14}	2^6
Maximum lag for ESS _{MCMC}	-	250	-	2000	5000
Bootstrap ESS _{MCMC}	-	*	-	72.7	47.3
Bridge ESS _{MCMC}	-	700.0	-	304.7	57.0
Bootstrap acceptance rate (%)	-	*	-	3.9	14.8
Bridge acceptance rate (%)	-	21.4	-	15.1	15.1
Configuration					
N ^o CPU threads	1	4	1	2	4
SSE instructions used?	No	No	No	No	Yes
GPU used?	No	No	No	Yes	No
Floating point precision	Double	Double	Double	Single	Double
Relative ESS threshold for resampling	0.5	0.5	0.5	0.5	0.5

* The bootstrap particle filter consistently degenerates on the FFR example, and no results could be obtained.

$\text{CAR}(z^{1:Z})$ is the *conditional acceptance rate* of $z^{1:Z}$ [34]:

$$\text{CAR}(z^{1:Z}) = \frac{1}{Z} \left(2 \sum_{i=1}^Z c^i - 1 \right),$$

where c^i is the sum of the i th smallest elements of $z^{1:Z}$. Higher values are favored for all three of the metrics (4.1)–(4.3).

The appropriate metric for comparison between methods depends on the motivation for estimating the normalizing constant. If the estimate itself is of interest, such as to compute evidence for model comparison, then the MSE-based metric (4.1) is appropriate. If the estimate is to be used as a weight in some importance sampling scheme, then the ESS-based metric (4.2) is most appropriate, as the ESS approximates the equivalent number of unweighted samples. If the estimate is instead to be used in some pseudomarginal MCMC scheme, such as

particle marginal Metropolis–Hastings (PMMH, Andrieu, Coucet, and Holenstein [1]), then the CAR-based metric (4.3) is most appropriate. CAR is an estimate of the long-term acceptance rate of a Metropolis chain that makes uniform proposals from Z states with posterior density proportional to the elements of $z^{1:Z}$ [34]. When exact likelihoods are computed, all elements of $z^{1:Z}$ are the same, and the CAR is one; in all other cases its difference from one represents the loss of using an estimated likelihood. Both ESS and CAR are sensitive to the high tail of $z^{1:Z}$, and reduce substantially in the presence of high outliers. They capture the dramatic loss in efficiency of importance and MCMC samplers in such circumstances. This is a particular risk when choosing weight functions for the bridge particle filter that are too tight. The MSE does not capture the implications of such outliers.

The OU, FFR, PD, and NPZD examples use a standard set of experiments for comparing the bootstrap and bridge particle filters. The SIR example does not, as the bootstrap particle filter could not be configured to work reliably on it for similar tests. For the toy OU and PD examples, parameters are fixed, and we simulate 16 data sets. For the FFR and NPZD examples, the data is fixed (a real-world data set), and we simulate 16 parameter sets from the prior for testing. We configure both the bootstrap and bridge particle filters with the number of particles set to each of $N \in \{2^5, 2^6, \dots, 2^{10}\}$. Using all combinations of the 16 data or parameter sets and six different settings for N gives 96 experiments in total. The three metrics are computed for each experiment, for a total of 288 comparisons on each example. The configurations for these tests are given in Table 1.

The FFR, SIR, and NPZD examples use real data sets. We perform parameter estimation in these scenarios using PMMH. We initialize two Markov chains, one using a bootstrap particle filter, and one using the bridge particle filter. Both are initialized to the same initial state, which has been obtained from a pilot run sufficiently long to have converged to the posterior distribution. Both use the same proposal distribution, tuned by hand from the same pilot run. The chains are compared using acceptance rate and effective sample size. The acceptance rate is a suitable comparison because the proposal distribution is the same for both chains, and a higher acceptance rate indicates less variability in normalizing constant estimates. The effective sample size used is that given in Kass et al. [23, p. 99]. This is different than that defined above, as it is intended for assessing the output of MCMC rather than that of importance sampling. It is defined as

$$\text{ESS}_{\text{MCMC}}(\theta^{1:N_\theta}) = \frac{N_\theta}{1 + 2 \sum_{k=1}^{\infty} R_\theta(k)},$$

where N_θ is the length of the chain and $R_\theta(k)$ its lag- k autocorrelation. In practice, $R_\theta(k)$ must be estimated from the chain itself, and the infinite sum truncated at some finite k , after which $R_\theta(k)$ is assumed to be zero. When there is more than one parameter, $\text{ESS}_{\text{MCMC}}(\theta^{1:N_\theta})$ is computed for each separately, and the smallest value reported.

4.1. Ornstein–Uhlenbeck (OU) process. Consider the OU process satisfying the following Itô SDE:

$$(4.4) \quad dX(t) = (\theta_1 - \theta_2 X(t)) dt + \theta_3 dW(t),$$




Figure 1. Comparison of particles generated by (left) the bootstrap particle filter and (right) the bridge particle filter for the OU example. The point at $(1, 0.15)$ in each plot indicates the observation. Solid vertical lines indicate times at which resampling is triggered. By introducing weighting and resampling steps at intermediate times, the bridge particle filter guides particles toward the observation.




Figure 2. Metrics for the OU example, comparing the bootstrap and bridge particle filters. Ninety-six experiments are conducted, each a unique combination of one of 16 simulated data sets and a number of particles $N \in \{2^5, 2^6, \dots, 2^{10}\}$. The bootstrap particle filter is used to estimate the normalizing constant for each experiment 4096 times, and each of the three metrics computed from these. The same is done for the bridge particle filter. Each figure then compares the bootstrap particle filter (x -axis) to the bridge particle filter (y -axis) for one of the metrics. Each point represents one of the experiments, the area of the point proportional to the number of particles in that experiment. Higher is better for all metrics, so that points above the diagonal are scenarios where the bridge particle filter rates superior, while points below are where the bootstrap particle filter rates superior. Red points are those where the bridge particle filter outperforms the bootstrap by a factor of at least two, and blue points are where the bootstrap particle filter outperforms the bridge by a factor of at least two. Remaining points are shown in grey.

with parameters $\theta_1 = 0.0187$, $\theta_2 = 0.2610$, and $\theta_3 = 0.0224$, as obtained in Aït-Sahalia [3] and used in Sun, Lee, and Hoeting [44]. For step size Δt , the transition density is [44]

$$p(x(t + \Delta t) | x(t)) = \phi(\mu(\Delta t), \sigma^2(\Delta t)),$$

with

$$\begin{aligned}\mu(\Delta t) &= \frac{\theta_1}{\theta_2} + \left(x(t) - \frac{\theta_1}{\theta_2}\right) \exp(-\theta_2 \Delta t), \\ \sigma^2(\Delta t) &= \frac{\theta_3^2}{2\theta_2} (1 - \exp(-2\theta_2 \Delta t)).\end{aligned}$$

Because the transition density is known explicitly for all Δt , no approximation of it is required, and Algorithm 3.1 can be applied.

We first consider sampling between the initial value $x(0) = 0$ and final value $x(1) = 0.15$, applying both the bootstrap and the bridge particle filters. The filters are configured as in Table 1. The results are shown in Figure 1. Clearly, the bridge particle filter produces a more satisfying result, with the additional weighting and resampling steps guiding particles towards the final value. Note that the performance of the bootstrap particle filter can be made arbitrarily poor in this example by reducing the discretization time step, pushing the value of $x(1) = 0.15$ further into the tails of the last transition density. This sensitivity of the bootstrap particle filter to the discretization time step would seem undesirable. If the bridge particle filter is allowed to resample after each step, it is not so sensitive.

Next, we compare the normalizing constant estimates of the bootstrap and bridge particle filters using the three metrics introduced above. We generate 16 data sets, each constructed by simulating the model forward from $x(0) = 0$ for 100 time units and taking the state at times $1, 2, \dots, 100$. The number of particles is set variously to $N \in \{2^5, 2^6, \dots, 2^{10}\}$. Each unique pair of a data set and an N constitutes an experiment, for 96 experiments in total. The bootstrap and bridge particle filters are applied to each experiment 4096 times, each time producing an estimate of the normalizing constant. From these estimates, each of the three metrics is computed. For computing the MSE-based metric, the true normalizing constant is used, this being readily computed as the model is linear and Gaussian. Results are in Figure 2. From this we see that the bridge particle filter outperforms the bootstrap particle filter in the great majority of comparisons, very often substantially so.

4.2. Federal funds rate (FFR). We apply the same process model (4.4) to 25 years of United States Federal Funds Rate data,² monthly from January 1989 to December 2013, with an interest in parameter estimation. A similar study is conducted in Aït-Sahalia [3]. Algorithm 3.1 can be used again. We put prior distributions on parameters

$$\begin{aligned}\theta_1 &\sim \mathcal{U}(-1, 1), \\ \theta_2 &\sim \mathcal{U}(0, 1), \\ \theta_3 &\sim \mathcal{U}(0, 1),\end{aligned}$$

where $\mathcal{U}(a, b)$ denotes a uniform distribution on the interval $[a, b]$.




Figure 3. Comparison of particles generated by (left) the bootstrap particle filter and (right) the bridge particle filter for the FFR example, over a short interval of time, with a particular parameter setting. Solid vertical lines indicate times at which resampling is triggered.




Figure 4. Metrics for the FFR example, comparing the bootstrap and bridge particle filters. See Figure 2 caption for explanation.




Figure 5. Posterior distributions over the state variable (first year only) and parameters for the FFR example, obtained with PMMH using the bridge particle filter. For the state variable, the shaded region indicates the 95% credibility interval at each time, and the middle line the median. Points mark observations. For parameters, a normalized histogram is given.

The first comparison is of filter behavior. Figure 3 compares the bootstrap and bridge particle filters over a short length of time. Observe the more satisfying behavior of the bridge particle filter as a result of the additional weighting and resampling steps.

The second comparison is of the normalizing constant estimates of the bootstrap and bridge particle filters using the three metrics introduced above. We simulate 16 parameter sets from the prior distribution. The number of particles is set variously to $N \in \{2^5, 2^6, \dots, 2^{10}\}$. Each unique pair of a parameter set and an N constitutes an experiment, for 96 experiments in total. The bootstrap and bridge particle filters are applied to each experiment 4096 times, each time producing an estimate of the normalizing constant. Each of the three metrics is computed from these estimates, for 288 comparisons in total. For computing the MSE-based metric, the true normalizing constant is used, this being readily computed as the model is linear and Gaussian. Results are in Figure 4.

The second comparison is to perform parameter estimation using a PMMH sampler. Two Markov chains are initialized from the same initial state, obtained from a pilot run that is considered to have converged to the posterior distribution. The same proposal distribution is used for both chains. Each chain is configured as in Table 1. The posterior distribution for the chain using the bridge particle filter is given in Figure 5, and its acceptance rate and effective sample size in Table 1. We have been unable to configure the bootstrap particle filter to work in this example, however. This is explained by the posterior of θ_3 being concentrated on very small values around 0.0017 to 0.0021 (see Figure 5). This results in a process with very narrow diffusivity, for which, it would seem, at least some observations become outliers with respect to the prior. These cannot be tracked by the bootstrap particle filter with a computationally feasible number of particles. The bridge particle filter can be tracked, even at this low θ_3 , because of the additional weighting and resampling steps.

4.3. Periodic drift (PD) process. Consider the diffusion process satisfying the following Itô SDE with nonlinear drift function, introduced in Beskos et al. [7] and studied further in Lin, Chen, and Mykland [27]:

$$(4.5) \quad dX(t) = \sin(X(t) - \theta) dt + dW(t).$$

We fix $\theta = \pi$ and $x(0) = 0$, and discretize using an Euler–Maruyama discretization at a time step of 0.075.

Algorithm 3.2 can be used. A Gaussian process does not capture the dynamics of this process well, as it is multimodal. We instead propose a parametric weight function

$$(4.6) \quad q(x_n | x_k) := \frac{1}{z} (\cos(x_n - \hat{x}_k) + 1 + \epsilon) \exp\left(-\frac{(x_n - \hat{x}_k)^2}{2\sigma^2(t_n - t_k)}\right),$$

where \hat{x}_k is x_k rounded to the nearest multiple of 2π , ϵ is a small positive value meant to prevent the density from being zero at the cosine troughs, and the normalizing constant z is

$$z = \sqrt{2\pi\sigma^2(t_n - t_k)} \left(\exp\left(-\frac{1}{2}\sigma^2(t_n - t_k)\right) + 1 + \epsilon \right).$$

²Obtained from <http://www.federalreserve.gov/releases/h15/data.htm>.




Figure 6. Weight function used for the PD example, (left) histogram of simulated paths used to fit the weight function, (center) the weight function (4.6) with maximum likelihood parameter estimates $\epsilon = 0.0259$ and $\sigma^2 = 0.3238$, and (right) the same weight function taken to the power 0.25 (used in the experiments). In all three cases the color scheme is scaled between zero and the greatest value in each vertical section.




Figure 7. Diffusion bridge samples for the PD example with the data set of Lin, Chen, and Mykland [27], $x(0) = 0$, $x(30) = 1.49$, $x(60) = -5.91$, $x(90) = -1.17$, using a (left) bootstrap particle filter, and (right) bridge particle filter. Solid vertical lines indicate times at which resampling is triggered.




Figure 8. Metrics for the PD example, comparing the bootstrap and bridge particle filters. See Figure 2 caption for an explanation.

To obtain the values of the parameters ϵ and σ^2 , we perform a maximum likelihood estimation using the Nelder–Mead method and a data set obtained by simulating 10000 paths from (4.5) for 30 units of time. It is important that this function is not too tight, or we risk high normalizing constant estimates that will dramatically reduce ESS and CAR. For this reason we take the final function to the power 0.25. The result is given in Figure 6.

As an initial test we simulate diffusion bridges conditioned on the data set in Lin, Chen, and Mykland [27], that is, $x(30) = 1.49$, $x(60) = -5.91$, $x(90) = -1.17$. Both the bootstrap and bridge particle filters are then applied. The number of particles is set to $N = 128$, with the bridge particle filter applying intermediate weighting and resampling at time steps of 1. The results are given in Figure 7. This shows that additional resampling is indeed triggered in the bridge particle filter on approaching each observation.

We next generate 16 data sets, each constructed by simulating the model forward for 3000 time units and taking the state at times $0, 30, \dots, 3000$. The number of particles is set variously to $N \in \{2^5, 2^6, \dots, 2^{10}\}$. Each unique pair of a data set and N constitutes an experiment, for 96 experiments in total. The bootstrap and bridge particle filters are applied to each experiment 4096 times to produce normalizing constant estimates, and each of the three metrics computed from these. For computing the MSE-based metric, a bootstrap particle filter using $N = 2^{20}$ is used to compute the “exact” log-likelihood for each parameter set. Results are in Figure 8.

4.4. Epidemiological SIR model. Consider an SIR (susceptible/infectious/recovered) model of an epidemic [24], where $S(t)$ gives the number of susceptible individuals in a population over time, $I(t)$ the number of infectious individuals, and $R(t)$ the number of recovered individuals, parameterized by an infection rate β and recovery rate ν :

$$\begin{aligned}\frac{dS(t)}{dt} &= -\beta S(t)I(t), \\ \frac{dI(t)}{dt} &= \beta S(t)I(t) - \nu I(t), \\ \frac{dR(t)}{dt} &= \nu I(t).\end{aligned}$$

We introduce stochasticity into the system by allowing the original parameters β and ν to vary in time [30, 15], following the Itô SDEs:

$$\begin{aligned}dS(t) &= -\beta(t)S(t)I(t) dt, \\ dI(t) &= (\beta(t)S(t)I(t) - \nu(t)I(t)) dt, \\ dR(t) &= \nu(t)I(t) dt, \\ d\log \beta(t) &= (\theta_{\beta,1} - \theta_{\beta,2} \log \beta(t)) dt + \theta_{\beta,3} dW_{\beta}(t), \\ d\log \nu(t) &= (\theta_{\nu,1} - \theta_{\nu,2} \log \nu(t)) dt + \theta_{\nu,3} dW_{\nu}(t).\end{aligned}$$

Note that $dS(t) + dI(t) + dR(t) = 0$, so that total population is conserved, and that $\beta(t)$ and $\nu(t)$ are always positive. For numerical simulation, the SDEs are converted to ODEs with a

discrete-time noise innovation:

$$\begin{aligned}\frac{dS(t)}{dt} &= -\beta(t)S(t)I(t), \\ \frac{dI(t)}{dt} &= \beta(t)S(t)I(t) - \nu(t)I(t), \\ \frac{dR(t)}{dt} &= \nu(t)I(t), \\ \frac{d \log \beta(t)}{dt} &= \theta_{\beta,1} - \theta_{\beta,2} \log \beta(t) + \theta_{\beta,3} \frac{\Delta W_\beta}{\Delta t}, \\ \frac{d \log \nu(t)}{dt} &= \theta_{\nu,1} - \theta_{\nu,2} \log \nu(t) + \theta_{\nu,3} \frac{\Delta W_\nu}{\Delta t},\end{aligned}$$

where each noise term $\Delta W \sim \mathcal{N}(0, \Delta t)$ is an increment of the Wiener process over a time step of size Δt . These ODEs are then numerically integrated forward using a low-storage fourth-order Runge–Kutta, with embedded third-order solution for error control, denoted RK4(3)5[2R+]C [8].

For the purposes of sampling diffusion bridges from the model, we wish to establish an ϵ -ball around observations within which samples must fall, where ϵ is comparable to the discretization error in simulating the model forward. It will make matters unnecessarily difficult to attempt to be any more accurate than this. The RK4(3)5[2R+]C algorithm outputs an error estimate for each variable at each time step, denoted $\epsilon_S(t)$, $\epsilon_I(t)$, and $\epsilon_R(t)$, computed as the difference between its third- and fourth-order solutions. A decision must then be made, in light of these errors, whether to accept or reject the step. The particular implementation [32] in LibBi is based on the description of error control for the DOPRI5 method in Hairer, Nørsett, and Wanner [22]. It uses an error tolerance parameterized by δ_{abs} (an absolute tolerance) and δ_{rel} (a relative tolerance). We set $\delta_{\text{abs}} = 10^{-2}$ and $\delta_{\text{rel}} = 10^{-5}$. For $X = \{S, I, R\}$, these are used to scale the error

$$(4.7) \quad \frac{\epsilon_X(t)}{\delta_{\text{abs}} + \delta_{\text{rel}} |X(t)|}.$$

The mean of these scaled errors is required to be less than one for the step to be accepted. If the step is rejected, the step size is reduced for a new attempt.

We suggest that it is sensible to use an ϵ -interval on each observed variable that is commensurate with this discretization error. We also note that the error estimate is of the *local* error of a single step of the numerical integrator, not of the *cumulative* error, which will be greater. We should, therefore, consider this estimate conservative, and may choose to inflate ϵ accordingly.

Only I is observed in the data set used below. We introduce an indirect (albeit highly informative) observation $Y_I(t)$, with uniform observation model:

$$Y_I(t) \sim \mathcal{U}(I(t) - \epsilon(t), I(t) + \epsilon(t)).$$

We can set $\epsilon(t) = \delta_{\text{abs}} + \delta_{\text{rel}} |I(t)|$, using the same error threshold as in (4.7), but this leads to the peculiar situation where the interval is wider for larger $I(t)$, so that the model grants higher

likelihood to smaller $I(t)$. This seems undesirable, so we remove the relative component. The overall population in the data set to be used is 763, so that the maximum error threshold is $\delta_{\text{abs}} + \delta_{\text{rel}} \times 763 = 0.01763$. Rounding up, we choose to set a constant $\epsilon = 0.02$ for all t . With this in place, the errors in the observation, like those in numerical integration, are kept accurate to within a small fraction of an individual of the population. We do not claim that this selection of ϵ is optimal, only that anything much less is futile without also tightening the error tolerances on the numerical integrator.

An alternative interpretation of this is as an ABC rejection algorithm with distance function $\rho(x(t), y(t)) := |I(t) - Y_I(t)|$ and acceptance criterion $\rho(x(t), y(t)) \leq \epsilon$, with the choice of ϵ guided by the discretization error of the numerical integrator.

We use Algorithm 3.3 with weight functions derived from the Gaussian process approach described in section 3. The variance of the additional observation noise is set to $\zeta^2(t) = \epsilon^2$. The data set records an outbreak of Russian influenza in a boys boarding school in northern England in 1978 [2]. The data set is also studied in Ross, Pagendam, and Pollett [41].

We put prior distributions over the parameters

$$\begin{aligned}\theta_{\beta,1} &\sim \mathcal{U}(-100, 100), \\ \theta_{\beta,2} &\sim \Gamma(2, 1), \\ \theta_{\beta,3} &\sim \mathcal{U}(0, 100), \\ \theta_{\nu,1} &\sim \mathcal{U}(-100, 100), \\ \theta_{\nu,2} &\sim \Gamma(2, 1), \\ \theta_{\nu,3} &\sim \mathcal{U}(0, 100),\end{aligned}$$

and the initial values of state variables

$$\begin{aligned}S(0) &= 763 - Y_I(0), \\ I(0) &= Y_I(0), \\ R(0) &= 0, \\ \log \beta(0) &\sim \mathcal{N}\left(\frac{\theta_{1,\beta}}{\theta_{2,\beta}}, \frac{\theta_{3,\beta}^2}{2\theta_{2,\beta}}\right), \\ \log \nu(0) &\sim \mathcal{N}\left(\frac{\theta_{1,\nu}}{\theta_{2,\nu}}, \frac{\theta_{3,\nu}^2}{2\theta_{2,\nu}}\right).\end{aligned}$$

Note that for $\log \beta(0)$ and $\log \nu(0)$, the prior is the same as the stationary distribution.

In preliminary experiments, we find that the bootstrap particle filter degenerates frequently for many settings of the parameters, while the bridge particle filter is much more reliable. Consequently, a PMMH chain using the bridge particle filter is configured as in Table 1 and used for a pilot run. This pilot run is continued until it appears to have converged to the posterior distribution, and has drawn enough samples to fit an improved random-walk Gaussian proposal. Initialized from the last state of the pilot run, we find that the bootstrap particle filter can now work more reliably and so may be used for comparison. We run two chains, one using the bootstrap and the other the bridge particle filter, both initialized from

the last state of the pilot run, and configured as in Table 1. Their resulting acceptance rates and effective sample sizes are reported there also. The chain using the bridge particle filter clearly performs better, although is slower also, given the additional resampling steps. The posterior results from this chain are given in Figure 9.

4.5. Marine biogeochemical NPZD model. Marine biogeochemical models are an important means of assessing ecosystem health, especially in coastal environments. We adopt the NPZD model of Parslow et al. [35]. The model is described in detail in that work, and summarized in Murray, Jones, and Parslow [34]. As even the summary takes several pages, we give only the high level motivation here.

An NPZD model represents the interaction of nutrients (N), phytoplankton (P), zooplankton (Z), and detritus (D) in a body of water. Each variable represents a compartment of a closed system, its value representing the quantity of nitrogen contained in that compartment. The four variables interact via a differential system where

$$\frac{dN(t)}{dt} + \frac{dP(t)}{dt} + \frac{dZ(t)}{dt} + \frac{dD(t)}{dt} = 0$$

such that the total quantity of nitrogen (in a closed system) is conserved. The fluxes between compartments are nonlinear functions of nine stochastic autoregressive terms, which model various biological, chemical, and physical processes on a discrete-time daily time step. A basic loop is the absorption of nutrient by phytoplankton during growth, the grazing of phytoplankton by zooplankton, the death of zooplankton to produce detritus, and the remineralization of that detritus into nutrient. The fluxes between variables are not limited to these particular interactions, however.

The NPZD model is physically positioned somewhere in the open ocean, within the surface mixed layer. There, it is subjected to exogenous environmental forcings such as daily temperature and light availability, and is opened by a bottom boundary condition that permits a flux of nutrient from below.

The data set used is from the site of Ocean Station P in the north Pacific. This data set has been studied before in Matear [31] and Parslow et al. [35]. We take four years of data, 1971–1974, as in Parslow et al. [35]. Observations include dissolved inorganic nitrogen (Y_N), considered an observation of nutrient (N), and chlorophyll-a fluorescence (Y_{Chla}), an observation of chlorophyll-a ($Chla$), itself a state variable that is a function of phytoplankton (P) and available light.

The observation model is

$$\begin{aligned}\log Y_N(t) &\sim \mathcal{N}(\log N(t), 0.2), \\ \log Y_{Chla}(t) &\sim \mathcal{N}(\log Chla(t), 0.5).\end{aligned}$$

While this observation model may appear only weakly informative, it can become highly informative given the sparsity (in time) at which observations are available. While the state variables tend to vary on daily or weekly time scales, the largest gap in observations of Y_N is 136 days, and that of Y_{Chla} 101 days.

We use Algorithm 3.3, with weight functions derived from the Gaussian process approach described in section 3. Gaussian processes are fit to the logarithm of the observed time series of Y_N and Y_{Chla} .

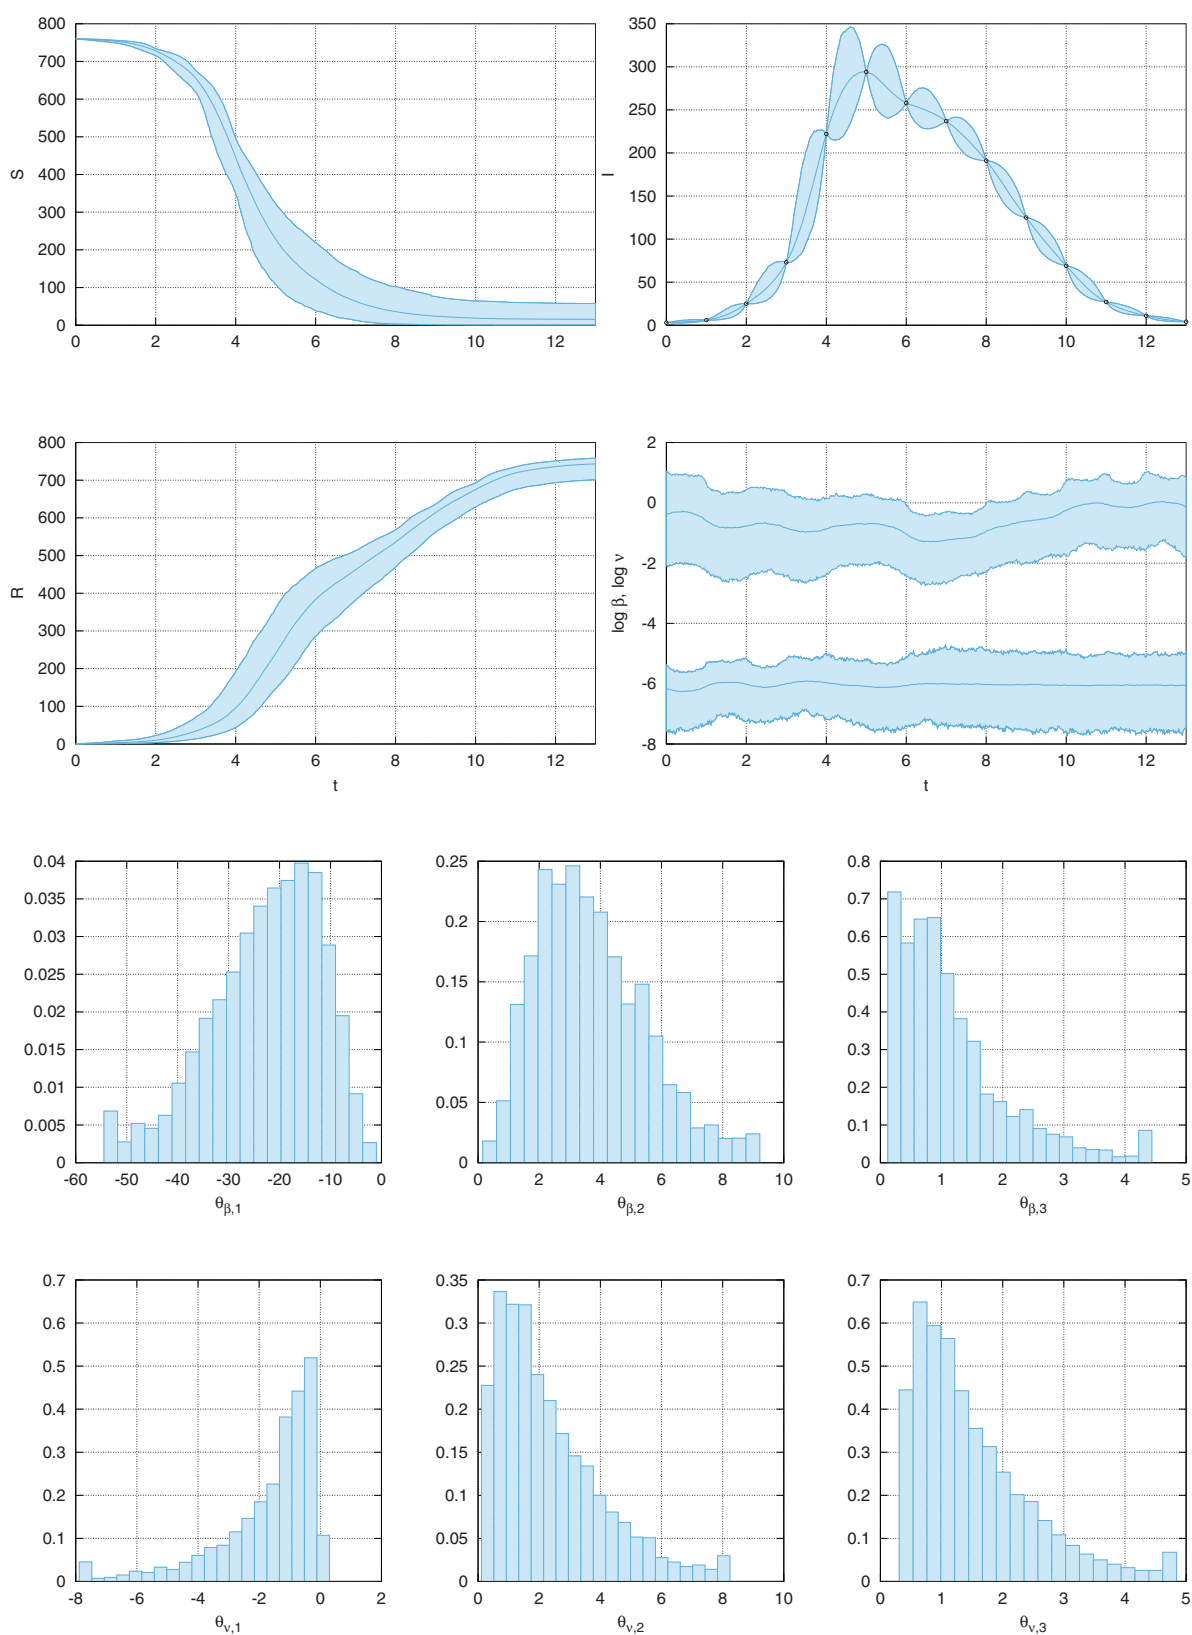


Figure 9. Posterior distributions over state variables and parameters for the SIR example, obtained with PMMH using the bridge particle filter. For state variables, shaded regions indicate 95% credibility intervals at each time, and middle lines the medians. Only I is observed, with observations marked. For parameters, a normalized histogram is given.




Figure 10. Metrics for the NPZD example, comparing the bootstrap and bridge particle filters. See Figure 2 caption for an explanation.

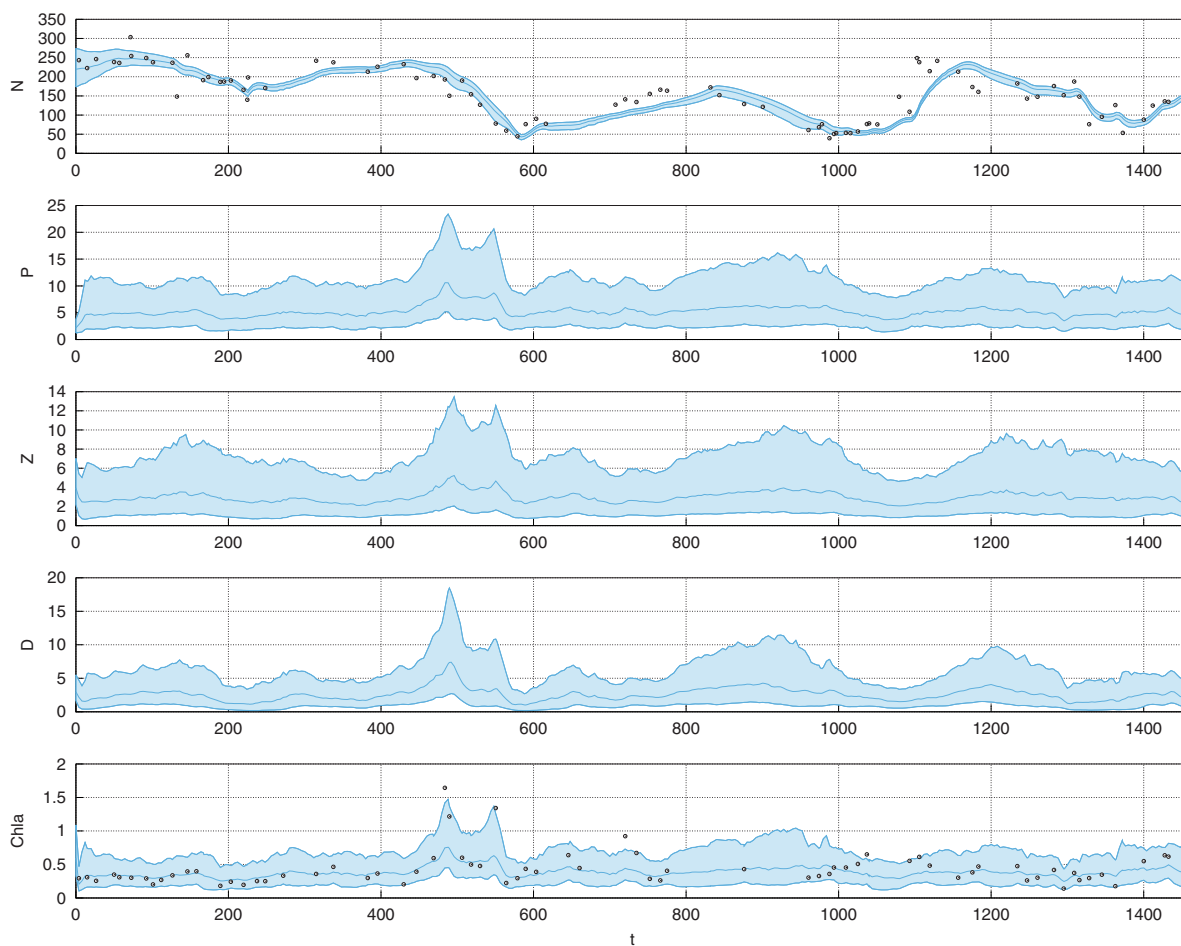


Figure 11. Posterior distributions over state variables for the NPZD example, obtained with PMMH using the bridge particle filter. Shaded regions indicate 95% credibility intervals at each time, and middle lines the medians. Both N and $Chla$ are observed, with observations marked; others are unobserved.

The first comparison is of the normalizing constant estimates of the bootstrap and bridge particle filters, using the three metrics introduced above. We simulate 16 parameter sets from the prior distribution. The number of particles is set variously to $N \in \{2^5, 2^6, \dots, 2^{10}\}$. Each unique pair of a parameter set and an N constitutes an experiment, for 96 experiments in total. The bootstrap and bridge particle filters are applied to each experiment 1024 times, each time producing an estimate of the normalizing constant. Each of the three metrics is computed from these estimates, for 288 comparisons in total. For computing the MSE-based metric, a bootstrap particle filter using $N = 2^{20}$ is used to compute the “exact” log-likelihood for each parameter set. Results are in Figure 10.

The second comparison is to perform parameter estimation using a PMMH sampler. Two Markov chains are initialized from the same initial state, obtained from a pilot run that is assessed to have converged to the posterior distribution. The same proposal distribution is used for both chains. They are otherwise configured as in Table 1, where their acceptance rates and effective sample sizes are also reported. The posterior distribution for the chain using the bridge particle filter is given in Figure 11.

5. Discussion. Of the four examples where the bootstrap and bridge particle filters are compared on metrics (OU, FFR, PD, NPZD), the bridge particle filter is consistently superior on the MSE metric, and at least as good on the ESS and CAR metrics. ESS is much better in the OU example, but similar on the others. CAR is much better in the OU example, modestly better in the FFR and PD examples, and similar in the NPZD example. This is evident in Figures 2, 4, 8, and 10. Recall that these metrics are already adjusted for execution time, which is typically longer for the bridge particle filter due to additional weighting and resampling steps.

Of the three examples where a comparison of PMMH performance using a real data set was attempted (FFR, SIR, and NPZD), we were unable to find a working configuration for the bootstrap particle filter for the FFR example, and had difficulties configuring it for the SIR example due to frequent degeneracy. The bridge particle filter worked reliably in both cases, however. In the SIR and NPZD examples, the bridge particle filter outperforms the bootstrap on both ESS and acceptance rate (see Table 1). This suggests that the bridge particle filter can work well in cases where the bootstrap particle filter does not.

We can report that the process of pilot runs—for tuning the number of particles and proposal distribution—was less unpleasant than usual when using the bridge particle filter. For the SIR example, the bootstrap particle filter could not be made to work for such pilot runs, so that the bridge particle filter was required for this purpose. While anecdotal, this experience does affirm that the additional weighting and resampling steps are useful, and may compensate for a poor setting of parameters, or poorly fitting model.

There are a number of areas where care is needed in configuring the bridge particle filter:

1. The additional resampling steps can decrease performance by introducing additional variance in the normalizing constant estimate. The use of an adaptive resampling trigger (such as the ESS used here) mitigates this. In the empirical results of this work, the additional resampling steps appear to have a net benefit, or in the worst cases, do no harm.
2. The weight functions used may be too tight, so that particles are selected too aggres-

sively at intermediate resamplings. For the PD, SIR, and NPZD examples we have taken the weight function to the power $1/4$ as a precaution. This seems sufficient for the examples here, but a more rigorous approach might be to use heavier-tailed weight functions (e.g., a Student t).

Finally, we have used a schedule of equispaced times for the additional weighting and resampling steps. This may be wasteful of computing resources. Alternative schedules may be superior, such as a geometric series of decreasing interval length, so that resampling is more frequent on approaching the observation. The schedule may even be adapted. We have found these ideas unnecessary to pursue for the examples here, however, and so leave them to future work.

6. Conclusion. This paper has presented three related SMC methods for handling state-space models with highly informative observations, including the special case of sampling bridges between fixed initial and final values. We have referred to them collectively as bridge particle filters. These bridge particle filters appear to improve substantially on the bootstrap particle filter in terms of the MSE of normalizing constant estimates, and more modestly on other metrics. For two applications, we have been able to apply a PMMH sampler using bridge particle filters when we either could not do so, or had difficulty doing so, with a bootstrap particle filter, and we report anecdotally that these methods are quite straightforward to configure.

Supplementary material. All examples are available for download from the LibBi website (www.libbi.org).

REFERENCES

- [1] C. ANDRIEU, A. DOUCET, AND R. HOLENSTEIN, *Particle Markov chain Monte Carlo methods*, J. R. Stat. Soc. Ser. B Stat. Methodol., 72 (2010), pp. 269–342.
- [2] ANONYMOUS, *Influenza in a boarding school*, British Medical Journal, 1 (1978), p. 587.
- [3] Y. AÏT-SAHALIA, *Transition densities for interest rate and other nonlinear diffusions*, J. Finance, 54 (1999), pp. 1361–1395.
- [4] C. BAYER AND J. SCHOENMAKERS, *Simulation of forward-reverse stochastic representations for conditional diffusions*, Ann. Appl. Probab. 24 (2014), pp. 1994–2032.
- [5] M. A. BEAUMONT, W. ZHANG, AND D. J. BALDING, *Approximate Bayesian computation in population genetics*, Genetics, 162 (2002), pp. 2025–2035.
- [6] M. A. BEAUMONT, J.-M. CORNUET, J.-M. MARIN, AND C. P. ROBERT, *Adaptive approximate Bayesian computation*, Biometrika, 96 (2009), pp. 983–990.
- [7] A. BESKOS, O. PAPASPILIOPOULOS, G. ROBERTS, AND P. FEARNHEAD, *Exact and computationally efficient likelihood-based estimation for discretely observed diffusion processes*, J. R. Stat. Soc. Ser. B Stat. Methodol., 68 (2006), pp. 333–382.
- [8] M. H. CARPENTER AND C. A. KENNEDY, *Fourth-Order 2N-Storage Runge–Kutta Schemes*, Technical Report Technical Memorandum 109112, National Aeronautics and Space Administration, Hampton, VA, 1994.
- [9] N. CHOPIN, P. E. JACOB, AND O. PAPASPILIOPOULOS, *SMC²: An efficient algorithm for sequential analysis of state space models*, J. R. Stat. Soc. Ser. B. Stat. Methodol., 75 (2013), pp. 397–426.
- [10] J. CLARK, *The simulation of pinned diffusions*, in Proceedings of the 29th IEEE Conference on Decision and Control, IEEE, Piscataway, NJ, 1990, pp. 1418–1420.
- [11] P. DEL MORAL, *Feynman-Kac Formulae: Genealogical and Interacting Particle Systems with Applications*, Springer-Verlag, New York, 2004.

- [12] P. DEL MORAL AND J. GARNIER, *Genealogical particle analysis of rare events*, Ann. Appl. Probab., 15 (2005), pp. 2496–2534.
- [13] P. DEL MORAL, A. DOUCET, AND A. JASRA, *An adaptive sequential Monte Carlo method for approximate Bayesian computation*, Stat. Comput., 22 (2012), pp. 1009–1020.
- [14] B. DELYON AND Y. HU, *Simulation of conditioned diffusion and application to parameter estimation*, Stochastic Process. Appl., 116 (2006), pp. 1660–1675.
- [15] J. DUREAU, K. KALOGEROPOULOS, AND M. BAGUELIN, *Capturing the time-varying drivers of an epidemic using stochastic dynamical systems*, Biostatistics, 14 (2013), pp. 541–555.
- [16] G. B. DURHAM AND A. R. GALLANT, *Numerical techniques for maximum likelihood estimation of continuous-time diffusion processes*, J. Bus. Econom. Statist., 20 (2002), pp. 297–316.
- [17] O. ELERIAN, S. CHIB, AND N. SHEPHARD, *Likelihood inference for discretely observed nonlinear diffusions*, Econometrica, 69 (2001), pp. 959–993.
- [18] B. ERAKER, *MCMC analysis of diffusion models with application to finance*, J. Bus. Econom. Statist., 19 (2001), pp. 177–191.
- [19] P. FEARNHEAD, O. PAPASPILIOPOULOS, AND G. O. ROBERTS, *Particle filters for partially observed diffusions*, J. R. Stat. Soc. Ser. B Stat. Methodol., 70 (2008), pp. 755–777.
- [20] P. FEARNHEAD, O. PAPASPILIOPOULOS, G. O. ROBERTS, AND A. STUART, *Random-weight particle filtering of continuous time processes*, J. R. Stat. Soc. Ser. B Stat. Methodol., 72 (2010), pp. 497–512.
- [21] A. GOLIGHTLY AND D. J. WILKINSON, *Bayesian sequential inference for nonlinear multivariate diffusions*, Stat. Comput., 16 (2006), pp. 323–338.
- [22] E. HAIRER, S. NØRSETT, AND G. WANNER, *Solving Ordinary Differential Equations I: Nonstiff Problems*, 2nd ed., Springer-Verlag, Berlin, 1993.
- [23] R. E. KASS, B. P. CARLIN, A. GELMAN, AND R. M. NEAL, *Markov chain Monte Carlo in practice: A roundtable discussion*, Amer. Statist., 52 (1998), pp. 93–100.
- [24] W. O. KERMACK AND A. G. MCKENDRICK, *A contribution to the mathematical theory of epidemics*, Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences, 115 (1927), pp. 700–721.
- [25] P. E. KLOEDEN AND E. PLATEN, *Numerical Solution of Stochastic Differential Equations*, Springer-Verlag, Berlin, 1992.
- [26] A. LEE, *Towards Smooth Particle Filters for Likelihood Estimation with Multivariate Latent Variables*, Master's thesis, University of British Columbia, Vancouver, British Columbia, 2008.
- [27] M. LIN, R. CHEN, AND P. MYKLAND, *On generating Monte Carlo samples of continuous diffusion bridges*, J. Amer. Statist. Assoc., 105 (2010), pp. 820–838.
- [28] M. LIN, R. CHEN, AND J. S. LIU, *Lookahead strategies for sequential Monte Carlo*, Statist. Sci., 28 (2013), pp. 69–94.
- [29] J. S. LIU AND R. CHEN, *Blind deconvolution via sequential imputations*, J. Amer. Statist. Assoc., 90 (1995), pp. 567–576.
- [30] X. LIU AND P. STECHLINSKI, *Infectious disease models with time-varying parameters and general nonlinear incidence rate*, Appl. Math. Model., 36 (2012), pp. 1974–1994.
- [31] R. J. MATEAR, *Parameter optimization and analysis of ecosystem models using simulated annealing: A case study at station P*, J. Mar. Res., 53 (1995), pp. 571–607.
- [32] L. M. MURRAY, *GPU acceleration of Runge-Kutta integrators*, IEEE Trans. Parallel Distrib. Syst., 23 (2012), pp. 94–101.
- [33] L. M. MURRAY, *Bayesian state-space modelling on high-performance hardware using LibBi*, in review, <http://arxiv.org/abs/1306.3277>, 2013.
- [34] L. M. MURRAY, E. M. JONES, AND J. PARSLow, *On disturbance state-space models and the particle marginal Metropolis-Hastings sampler*, SIAM/ASA J. Uncertain. Quantif., 1 (2013), pp. 494–521.
- [35] J. PARSLow, N. CRESSIE, E. P. CAMPBELL, E. JONES, AND L. M. MURRAY, *Bayesian learning and predictability in a stochastic nonlinear dynamical model*, Ecological Applications, 23 (2013), pp. 679–698.
- [36] A. PEDERSEN, *A new approach to maximum likelihood estimation for stochastic differential equations based on discrete observations*, Scand. J. Statist., 22 (1995), pp. 55–71.
- [37] G. W. PETERS, Y. FAN, AND S. A. SISSON, *On sequential Monte Carlo, partial rejection control and approximate Bayesian computation*, Stat. Comput., 22 (2012), pp. 1209–1222.

- [38] M. PITT AND N. SHEPHARD, *Filtering via simulation: Auxiliary particle filters*, J. Amer. Statist. Assoc., 94 (1999), pp. 590–599.
- [39] M. K. PITT, *Smooth Particle Filters for Likelihood Evaluation and Maximisation*, Technical Report 651, Department of Economics, The University of Warwick, Warwick, UK, 2002.
- [40] G. O. ROBERTS AND O. STRAMER, *On inference for partially observed nonlinear diffusion models using the Metropolis–Hastings algorithm*, Biometrika, 88 (2001), pp. 603–621.
- [41] J. V. ROSS, D. E. PAGENDAM, AND P. K. POLLETT, *On parameter estimation in population models II: Multi-dimensional processes and transient dynamics*, Theoret. Population Biol., 75 (2009), pp. 123–132.
- [42] M. SCHAUER, F. VAN DER MEULEN, AND H. VAN ZANTEN, *Guided Proposals for Simulating Multi-dimensional Diffusion Bridges*, <http://arxiv.org/abs/1311.3606>, 2013.
- [43] S. A. SISSON, Y. FAN, AND M. M. TANAKA, *Sequential Monte Carlo without likelihoods*, Proc. Natl. Acad. Sci. USA, 104 (2007), pp. 1760–1765.
- [44] L. SUN, C. LEE, AND J. A. HOETING, *Penalized Importance Sampling for Parameter Estimation in Stochastic Differential Equations*, <http://arxiv.org/abs/1305.4390>, 2013.
- [45] N. G. VAN KAMPEN, *Stochastic Processes in Physics and Chemistry*, North–Holland, Amsterdam, New York, 2007.